arXiv:1702.02241v1 [math.OC] 8 Feb 2017

Adapting Regularized Low Rank Models for Parallel Architectures

Derek Driggs*, Stephen Becker*, Aleksandr Aravkin†

February 9, 2017

Abstract

We introduce two new methods to parallelize low rank recovery models in order to take advantage of GPU, multiple CPU, and hybridized architectures. Using Burer-Monteiro splitting and marginalization, we develop a smooth, non-convex formulation of regularized low rank recovery models that can be solved with first-order optimizers. Using L-BFGS on the GPU, the resulting non-convex programs offer enormous speedup over existing methods.

Our second development is a communication-avoiding implementation of the randomized singular value decomposition (rSVD) that ameliorates the heavy communication accompanying tall-skinny matrix structures that arise in intermediate calculations of low rank recovery problems. We demonstrate, using synthetic data, surveillance video, and data from fMRI brain scans, that both algorithms offer significant speedup over traditional low rank recovery models.

1 Introduction

Low rank matrix decompositions are an effective tool for large-scale data analytics. Background subtraction, facial recognition, document indexing, and collaborative filtering all use low rank matrix recovery algorithms (Candès et al., 2011). However, a low rank structure is only a part of the story — signals often exhibit additional structure. *Regularized* low rank recovery models seek to decompose a matrix $X \in \mathbb{R}^{m\times n}$ into the sum of a low rank component $L \in \mathbb{R}^{m\times n}$ and another well-structured matrix $S \in \mathbb{R}^{m\times n}$, where S could be sparse, clustered, non-negative, or have some other useful attribute (see e.g. Udell et al. (2015)).

Searching for sparse and low rank models gives rise to combinatorial formulations, which are intractable. Fortunately, relaxations often work well, and are amenable to modern optimization techniques. Consider finding an approximate decomposition of a data matrix X into a sum $X = L + S$, where S is sparse and L low rank. We are given partial linear measurements $b = \mathcal{A}(X)$. We can formulate the problem

$$\min_{L,S} \quad \text{rank}(L) + \text{card}(S), \quad \text{subject to:} \quad \frac{1}{2}\|\mathcal{A}(L+S) - b\|_2^2 \leq \epsilon, \tag{1}$$

where the cardinality function $\text{card}(\cdot)$ counts the number of nonzero entries of S, and the parameter ϵ accounts for noise or model inconsistencies. Rank is analogous to sparsity, since it is precisely the $\text{card}(\cdot)$ function applied to the singular values of the argument. Problem (1) is NP-hard (Natarajan, 1995) and intractable for large problem sizes. However, both terms admit nonsmooth convex relaxations. $\text{card}(\cdot)$ can be replaced with the ℓ_1-norm, and the rank function can be replaced with the nuclear norm, $\|\cdot\|_*$, which is equal to the sum of the singular values of the argument:

$$\min_{L,S} \quad \|L\|_* + \|S\|_1, \quad \text{subject to:} \quad \frac{1}{2}\|\mathcal{A}(L+S) - b\|_2^2 \leq \epsilon \tag{2}$$

When the parameter ϵ is set to zero, (2) is called *robust PCA* (RPCA), and can recover the lowest-rank L and the sparsest S under mild conditions. Setting $\epsilon > 0$ in (2) is more suitable for most applications, since it allows for noise in the measurements; the problem is then called *stable principal component pursuit* (SPCP) (Wright et al., 2009; Zhou et al., 2010). SPCP stably recovers L and S with error bounded by ϵ (Zhou et al., 2010). Both RPCA and SPCP are convex problems, much easier to solve than the original problem (1). While (2) is an important problem, it is just one example of regularized low rank models.

*University of Colorado, Boulder CO 80309 USA
†University of Washington, Seattle, WA 98195 USA

Problem class. We are interested in general regularized low rank approaches:

$$\min_{L,S} \quad \|L\|_* + \mathcal{L}(b, \mathcal{A}(L+S)) + r(S), \tag{3}$$

where $\mathcal{L}$ is a loss function and r is a convex regularizer. $\mathcal{L}$, which can be *infinite-valued*, measures how well $L+S$ agrees with measured values of X. In RPCA, $\mathcal{L}(v) = \delta_b(v)$ is the indicator function ensuring we match observed data, while in SPCP, $\mathcal{L}(v) = \delta_{2\sqrt{\epsilon}\mathbb{B}}(v-b)$ ensures we are close to b. $\mathcal{L}$ can also be finite valued — for example, a modified SPCP is given by solving

$$\min_{L,S} \quad \lambda_L \|L\|_* + \frac{1}{2}\|\mathcal{A}(L+S) - b\|_2^2 + \lambda_S \|S\|_1, \tag{4}$$

where λ_L, λ_S are tuning parameters, and $\mathcal{L}(\cdot) = \frac{1}{2}\|\cdot\|_2^2$.

The regularizer $r(S)$ is a convex penalty that promotes desired structure of S. The ℓ_1 norm $r(S) = \|S\|_1$ promotes sparsity in S; other important examples include the ordered-weighted ℓ_1 norm (OWL norm) (Zeng & Figueiredo (2014); Davis (2015); Bogdan et al. (2015), and the elastic net (Zou & Hastie, 2005), which both enforce sparsely correlated observations (clusters) in S. Other convex regularizers promote known group-sparse and hierarchical structures in S; see Bach et al. (2012) for a survey.

Marginalization in S. In all of the applications we consider, $r(S)$ is simple enough that we can efficiently minimize (3) in S for a given L. In the motivating SPCP example (4), we have

$$\min_S \left\{ \lambda_L \|L\|_* + \frac{1}{2}\|\mathcal{A}(L+S) - b\|_2^2 + \lambda_S \|S\|_1 \right\} = \lambda_L \|L\|_* + \rho_{\lambda_S}(\mathcal{A}(L) - b),$$

where ρ_{λ_S} is the Huber function (for a detailed derivation, see e.g. Aravkin et al. (2016)). This motivates the definition

$$\varphi : L \mapsto \min_S \left\{ \mathcal{L}(\mathcal{A}(L+S) - b) + \lambda_S \cdot r(S) \right\}, \tag{5}$$

where φ is smooth (see section 3.2) with value and derivative either explicitly available as in the case of SPCP, or efficiently computable. After marginalizing out S, (3) becomes

$$\min_L \ \lambda_L \|L\|_* + \varphi(L). \tag{6}$$

Proximal gradient methods are applicable, but require computing the expensive singular value decomposition (SVD) at each iteration. To avoid this bottleneck, we develop and compare two alternatives for accelerating the method: parallelizing SVD computations using the tall-skinny QR algorithm, and non-convex factorized representations that avoid the SVD entirely. The smooth structure of φ allows a computable *optimality certificate* for the non-convex case, which is developed and explored in section 3.

Accelerating SVD computation. The SVD is notoriously costly to compute, and its high communication costs make it difficult to parallelize. Making matters worse, in many applications the data matrix has many more rows than columns, and these "tall-skinny" matrices make the communication costs of a parallel SVD much higher.

To accelerate the SVD step, the randomized SVD (rSVD) (Halko et al., 2011) is often used to efficiently compute a partial SVD (see e.g. Aravkin et al. (2014a)). While the rSVD algorithm requires fewer operations than traditional SVDs, intermediate steps of the algorithm exacerbate the tall-skinny problem, which makes the rSVD even more difficult to parallelize. In order to use parallelized architectures such as the GPU to accelerate regularized low rank recovery, the communication cost of the SVD step must be reduced. To do this, we use the tall-skinny QR algorithm (TSQR) of Demmel et al. (2012). Anderson et al. (2010) explored using TSQR to accelerate the SVD step of RPCA for tall-skinny matrices, but we address an inherent communication problem in the rSVD algorithm that arises even for square matrices.

Non-convex reformulation. To avoid the SVD entirely, we consider non-convex reformulations for regularized low rank recovery models, using the factorization $L = UV^T$, where $U \in \mathbb{R}^{m\times k}$, $V \in \mathbb{R}^{n\times k}$, and $k \ll m, n$. This factorization imposes the rank-constraint $\text{rank}(L) \leq k$. In addition, we can further penalize rank using the representation (see e.g. Srebro et al. (2005); Recht et al. (2010); Bach et al. (2008))

$$\|L\|_* \equiv \inf_{L=UV^T} \frac{1}{2}(\|U\|_F^2 + \|V\|_F^2). \tag{7}$$

This equivalence only holds if the ranks of U and V are high enough, and makes it possible to maintain a low rank $L = UV^T$ without requiring SVDs. The non-convex analogue to (6) is given by

$$\underset{U,V}{\arg\min} \quad \varphi(UV^T) + \frac{\lambda_L}{2}(\|U\|_F^2 + \|V\|_F^2). \tag{8}$$

Roadmap. We start with a survey of related work in section 2. We then consider two theoretical issues related to the smooth reformulation (8) in section 3 . First, using the factorization $L = UV^T$ gives a non-convex program that could potentially create new stationary points and local minima. Second, it is important to understand the properties of the map φ obtained by marginalizing over S in the context of specific optimization algorithms. In section 4, we study heuristic initialization methods, which are important for non-convex formulations. We present detailed numerical studies, comparing the non-convex method with competing alternatives, in section 5. The convex formulation (6) and TSQR method are presented in section 6, along with additional numerical experiments, illustrating the parallelism of TSQR and comparing it to the non-convex method. We end with final thoughts in section 7.

2 Prior Work and Contributions

Several authors have used alternating minimization to solve non-convex, factorized matrix completion models (Hardt, 2013; Keshavan, 2012; Keshavan & Montanari, 2010; Keshavan et al., 2009, 2010; Gunasekar et al., 2013; Jain et al., 2013). This line of research is supported by the existence of conditions that guarantee fast convergence to a globally optimal solution (Hardt, 2013; Keshavan & Montanari, 2010; Keshavan et al., 2010, 2009; Gunasekar et al., 2013), sometimes at a linear rate (Jain et al., 2013; Udell et al., 2015). These models are a special class of (8), when $r(S) = 0$.

In Udell et al. (2015), the authors review the use of alternating minimization to solve regularized PCA problems of the following form:

$$\min_{U,V} \quad \mathcal{L}(UV^T, X) + r(U) + \tilde{r}(V). \tag{9}$$

Matrix completion models and RPCA can adopt the split-form in (9), but SPCP and other regularized models cannot. Our analysis includes more general regularizers and offers an alternative to alternating minimization.

Shen et al. (2012) develop a split RPCA program and solve it using alternating minimization. While their technique can be an order of magnitude faster than competitive convex solvers (such as the inexact augmented Lagrangian method (IALM) from Lin et al. (2011)), its performance suffers when the magnitude of the sparse component is large compared to the magnitude of the low rank component.

Tu et al. (2016) follow techniques introduced in Candès et al. (2015) and consider provable recovery of a low rank solution to a system of linear equations using a non-convex, factorized program. In **?**, the authors extend these ideas to investigate a non-convex formulation of robust PCA that can be solved with gradient descent.

Another non-convex approach to RPCA is considered in Netrapalli et al. (2014). In contrast to the research discussed previously, their method does not use a low rank factorization of L. Instead, these authors develop an algorithm of alternating projections, where each iteration sequentially projects L onto a set of low rank matrices and thresholds the entries of S. Their technique also comes with recovery and performance guarantees.

We develop a general non-convex formulation for regularized low rank recovery models that can be solved with general first-order optimization algorithms. We address some problems associated with local minima

and spurious stationary points that accompany non-convexity, by showing that all local minima of our model have the same objective value, and by providing a certificate to prove that a given solution is not a spurious stationary point. We also show that our method is particularly well-suited for the GPU.

We also develop a communication-efficient solver for convex models by limiting the communication costs in the rSVD, using TSQR to replace the orthogonalization step. Anderson et al. (2010) briefly discusses using TSQR to accelerate the full SVD step in RPCA, but our rSVD algorithm offers greater speedup and applies to general matrices, not just tall and skinny matrices. In Martinsson & Voronin (2015), the authors propose a blocked algorithm to accelerate the rSVD. Their blocking scheme reduces communication costs in a way similar to TSQR, but TSQR also offers additional parallelism. Both of our approaches offer significant speedup over existing methods, as illustrated in section 5.

3 Theoretical Considerations

We discuss three issues related to the approach (8). First, we study the convexity, smoothness, and differentiability of the marginal function (6) in section 3.2. Next, moving from (6) to (8) by factorizing L, we transform a convex problem into a non-convex one. This creates two issues: an implicit rank constraint $\operatorname{rank}(L) \leq k$ that is not present in (8), and potential for local minima and spurious stationary points in the non-convex variant that do not correspond to a solution of the convex problem. We address these issues in section 3.3.

Section 3.4 introduces a computable certificate that can be used to check whether the non-convex model has converged to a global solution. Using the smoothness of φ, we derive a simple optimality certificate that can be tracked while optimizing (8) to identify either when the implicit rank of the factors is too small or when we have converged to a minimizer of (6).

3.1 Convexity of φ

The convexity of the marginal function (6) is a well-known result in convex analysis (Rockafellar & Wets, 2009, Proposition 2.22). We include a short, self-contained proof of this fact.

Lemma 1. *Let $f : \mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n} \to \mathbb{R}$ be a jointly convex function, and let $L \in \mathbb{R}^{m\times n}$. The function $g : \mathbb{R}^{m\times n} \to \mathbb{R} : L \mapsto \min_{S\in\mathbb{R}^{m\times n}} f(L, S)$ is convex.*

Proof. For all $L_1, L_2, S_1, S_2 \in \mathbb{R}^{m\times n}$ and $t \in (0, 1)$,

$$\begin{aligned} g(tL_1 + (1-t)L_2) &= \min_{S\in\mathbb{R}^{m\times n}} f(tL_1 + (1-t)L_2, S) \\ &\leq f(tL_1 + (1-t)L_2, tS_1 + (1-t)S_2) \\ &\leq tf(L_1, S_1) + (1-t)f(L_2, S_2). \end{aligned}$$

As this inequality holds for arbitrary S_1 and S_2, it holds for any minimizing $S^\star$:

$$\begin{aligned} g(tL_1 + (1-t)L_2) &\leq tf(L_1, S^\star) + (1-t)f(L_2, S^\star) \\ &= tg(L_1) + (1-t)g(L_2). \end{aligned}$$

□

It follows immediately that φ as defined in (5) is convex when $\mathcal{L}$ and r are convex.

3.2 Smoothness of φ

When $\mathcal{L}$ is the least-squares error and

$$\mathcal{A}(L + S) = \mathcal{A}_L(L) + \mathcal{A}_S(S),$$

with $\mathcal{A}_S$ an injective linear operator, Lipschitz differentiability of φ is immediate from the strong convexity of $\mathcal{L}$. Support for this claim is given, for example, in Bauschke & Combettes (2011), where the relevant proposition for the specific case where $\mathcal{A}_S$ is a reshaping operator is as follows:

Proposition ***(Bauschke & Combettes, 2011, Prop. 12.29):*** *Let r be a proper lower-semicontinuous convex function, and let ${}^{\mu}r(\mathcal{A}(L)-\mathrm{vec}(S)-b) = \inf_S \; r(S)+\frac{1}{2\mu}\|\mathcal{A}(L)-\mathrm{vec}(S)-b\|_2^2$ be the Moreau envelope of f with parameter μ. Then ${}^{\mu}r$ is differentiable and its gradient is μ^{-1}-Lipschitz continuous.*

More generally, $\varphi(L)$ is differentiable as long as $r(S)$ has a unique minimizer and the loss function $\mathcal{L}$ is convex and smooth. By (Rockafellar & Wets, 2009, Theorem 10.58), we have

$$\partial\varphi(\cdot) = \mathcal{A}^*\nabla\mathcal{L}\left(\mathcal{A}(\cdot) - b - \mathcal{A}(S^\star)\right), \tag{10}$$

where $S^\star$ is any minimizer of (4). The subdifferential is a singleton if and only if the minimizing $S^\star$ is unique, and a unique minimum is guaranteed if the sum $\mathcal{L}(\mathcal{A}(L+S)-b)+r(S)$ in (4) is strictly convex.

Using SPCP as an example, we have

$$\varphi(L) = \min_S \left\{ \frac{1}{2}\|\mathcal{A}(L)+\mathrm{vec}(S)-b\|_2^2 + \lambda_S\|S\|_1 \right\},$$

where sparse outliers S are fully observed. The objective is strongly convex in S, and has the unique minimizer

$$S^\star = \mathrm{prox}_{\lambda_S\|\cdot\|_1}(\mathcal{A}(L)-b).$$

Therefore, $\varphi(L)$ is smooth, with

$$\nabla\varphi(L) = \mathcal{A}^*\left(\mathcal{A}(L)+\mathrm{vec}(S^\star)-b\right).$$

Note we used (10) rather than the parametric form of $\varphi(L)$ to obtain its value and gradient. Using (7) to replace the nuclear norm and replacing the non-smooth regularizer with φ, we have created a smooth optimization problem amenable to first-order methods.

For minimizing smooth objectives, quasi-Newton methods are often faster than algorithms that require only first-order smoothness, such as gradient-descent. The convergence rates of quasi-Newton methods depend on second-order smoothness, which does not hold in problems such as SPCP, since e.g. the Huber function is only $\mathcal{C}^1$. However, empirical results (including those presented in section 5) suggest that these methods are still effective.

3.3 Rank and Local Minima

The factorized problem (8) is not equivalent to (6) because of the implicit rank constraint. However, we can show that (8) is equivalent to the rank-constrained problem

$$\min_L \quad \lambda_L\|L\|_* + \varphi(L) \quad \text{subject to:} \quad \mathrm{rank}(L) \le k. \tag{11}$$

This follows immediately from (7). In order for the split problem to recover the solution to (1), it must be initialized with a k larger than the rank of the minimizing L. However, a larger k slows computation, so it must be chosen with care. This issue is considered in more depth in section 4.

We want to be sure that any local minimum of (8) corresponds to a (global) minimum of (6). The following theorem combines ideas from Burer & Monteiro (2005) and Aravkin et al. (2014b) to show this holds provided that k is larger than the rank of the minimizer for (6).

Theorem 1. *Consider an optimization problem of the following form:*

$$\min_{X\succeq 0} \quad f(X), \quad \textit{such that} \quad \mathrm{rank}(X) \le k, \tag{12}$$

where $X \in \mathbb{R}^{n\times n}$ is a positive semidefinite real matrix, and f is a lower semi-continuous function mapping to $[-\infty,\infty]$ and has a non-empty domain over the set of positive semi-definite matrices. Using the change of variable $X = PP^T$, take $P \in \mathbb{R}^{n\times k}$, and consider the problem

$$\min_P \quad g(P) \stackrel{def}{=} f(PP^T). \tag{13}$$

Let $\bar{X} = \bar{P}\bar{P}^T$, where $\bar{X}$ is feasible for (12). Then $\bar{X}$ is a local minimizer of (12) if and only if $\bar{P}$ is a local minimizer of (13).

The proof of Theorem 1 is deferred to Appendix A. Using the SDP formulation of the nuclear norm presented in Fazel et al. (2001), our problem can be recast as a semi-definite program so that we can apply Theorem 1. Define

$$Z = \begin{bmatrix} U \\ V \end{bmatrix} \begin{bmatrix} U \\ V \end{bmatrix}^T = \begin{bmatrix} UU^T & L \\ L^T & VV^T \end{bmatrix}. \tag{14}$$

The matrix Z is positive semi-definite, and has form $Z = PP^T$, with $P = \begin{bmatrix} U \\ V \end{bmatrix}$. Let $\mathcal{R}(\cdot)$ be the function that extracts the upper-right block of a matrix (so that $\mathcal{R}(Z) = L$), and let

$$f(Z) = \frac{\lambda_L}{2} \operatorname{trace}(Z) + \varphi(\mathcal{R}(Z)). \tag{15}$$

Using equation (3), we now see that the rank-constrained problem is equivalent to

$$\min_{Z \succeq 0} \quad f(Z), \quad \text{such that} \quad \operatorname{rank}(Z) \leq k. \tag{16}$$

Applying Theorem 1, we can now be assured that $\overline{P}$ is a local minimizer to the split program if and only if $\overline{Z}$ is local minimizer of the original problem. This is equivalent to the statement that the point $(\overline{U}, \overline{V})$ is a local minimizer of the split problem if and only if $\overline{L} = \overline{U}\ \overline{V}^T$ is a local minimizer of the original, rank-constrained program.

3.4 A Certificate for Convergence

Although Theorem 1 asserts that the split problem and the rank-constrained problem have the same local minima, the rank-constrained problem is itself non-convex, so we have not guaranteed that every stationary point of the split problem solves the convex program. Using an approach building on work of Udell et al. (2015), we can develop a certificate to check whether a recovered solution to (6) corresponds to a spurious stationary point when φ is convex (in particular when $\mathcal{L}$ and r are convex). This technique can also be used on the convex formulation as a method to check the distance to optimality.

We base our certificate on the following convex analysis. Let F be any proper convex function, then Fermat's rule states that the set of minimizers of F are points L such that $0 \in \partial F(L)$, where $\partial(\cdot)$ denotes the subdifferential. One could take a candidate point L, evaluate $F(L)$, and if this is zero, conclude that L is the minimizer, but such expectations are unrealistic in floating point arithmetic, much less from an iterative method. Suppose instead we find

$$\mathcal{E} \in \partial F(L) \tag{17}$$

where $\mathcal{E}$ is small. Let $L^\star$ denote any minimizer of F, then from the definition of a subgradient, $F(L^\star) \geq F(L) + \langle \mathcal{E}, L^\star - L \rangle$. Then

$$F(L) - F(L^\star) \leq \langle \mathcal{E}, L - L^\star \rangle \leq \|\mathcal{E}\|_p \cdot \|L - L^\star\|_d \tag{18}$$

for any pair of primal and dual norms $\|\cdot\|_p$ and $\|\cdot\|_d$. We discuss bounding $\|\mathcal{E}\|_p$ and $\|L - L^\star\|_d$ in the next two subsections.

Finding an approximate zero $\mathcal{E}$ The key to the certificate is analyzing the marginalized problem (6) in terms of L, rather than the problem (3) in terms of (L, S). The point $L = UV^T$ is the optimal solution of (6) if and only if

$$0 \in \partial\big(\|L\|_* + \varphi(L)\big). \tag{19}$$

We have set $\lambda_L = 1$ for convenience, or alternatively one can absorb λ_L^{-1} into φ. Since both the nuclear norm and φ are proper lower semi-continuous and the intersection of the interior of their domains is non-empty (in particular they are finite valued), then by Bauschke & Combettes (2011, Cor. 16.38), we have that

$$0 \in \partial\big(\|L\|_* + \varphi(L)\big) \Leftrightarrow 0 \in \big(\partial\|L\|_*\big) + \big(\partial\varphi(L)\big).$$

Both of these subdifferentials are computable. Let

$$L = \widetilde{U}\Sigma\widetilde{V}^T = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} \Sigma_1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix} \tag{20}$$

be the (full) SVD of L, where $U_1 \in \mathbb{R}^{m\times r}$ and $V_1 \in \mathbb{R}^{n\times r}$. The subdifferentials of the nuclear norm at L comprises matrices of the form $X = U_1V_1^T + W$, where U_1, V_1 contain the left and right singular vectors of L that correspond to non-zero singular values,[1] and W satisfies the conditions $U_1^TW = 0$, $WV_1 = 0$ and $\|W\|_2 \leq 1$. Equivalently,

$$X \in \partial\|L\|_* \quad \text{if and only if} \quad \exists W' \in \mathbb{R}^{m-r\times n-r} \text{ s.t. } \widetilde{U}^TX\widetilde{V} = \begin{bmatrix} I & 0 \\ 0 & W' \end{bmatrix}, \quad \|W'\| \leq 1. \tag{21}$$

We can find an explicit matrix $D \in \partial\varphi(L)$ using (10), and it is unique under the smoothness conditions described in Section 3.2.

Altogether, we can guarantee a point $\mathcal{E} \in \partial\|L\|_* + \partial\varphi(L)$ such that

$$\begin{aligned} \|\mathcal{E}\|_F^2 &= \min_{X\in\partial\|L\|_*} \|X + D\|_F^2 \\ &= \min_{X\in\partial\|L\|_*} \|\widetilde{U}^T(X + D)\widetilde{V}\|_F^2 \\ &= \min_{\|W'\|\leq 1} \left\| \begin{bmatrix} I & 0 \\ 0 & W' \end{bmatrix} + \widetilde{U}^TD\widetilde{V} \right\|_F^2 \\ &= \|I - U_1^TDV_1\|_F^2 + \|U_1^TDV_2\|_F^2 + \|U_2^TDV_1\|_F^2 + \min_{\|W'\|\leq 1} \|U_2^TDV_2 - W\|_F^2. \end{aligned} \tag{22}$$

In the first equation, the minimum is achieved since the squared-norm is continuous and the set is compact, and (21) is used in the third equation. Each term of the fourth equation can be calculated explicitly, with the last term obtained via projection onto the unit spectral-norm ball, which requires the SVD of the small matrix $U_2^TDV_2$. This bound on $\|\mathcal{E}\|_F$ can then be used in (18).

Most terms above can be computed efficiently, in the sense that U_2 and V_2 never need to be explicitly computed (which is most important when r is small, since then U_2 and V_2 have $m - r$ and $n - r$ columns, respectively), and therefore the computation consists only of matrix multiplies, thin-QR factorizations, and an $r \times r$ SVD, leading to a complexity of $\mathcal{O}(rmn + r^2 \cdot (m + n) + r^3)$. For example, we compute

$$\|U_1^TDV_2\|_F^2 = \|U_1^TD\|_F^2 - \|U_1^TDV_1\|_F^2$$

and $\|U_2^TDV_1\|_F^2$ is computed analogously. The final term requires an unavoidable $(m - r) \times (n - r)$ SVD factorization, so for large m, n this is not cheap, and in this case we suggest only computing the certificate as a final check at the end rather than at every point to check convergence.

Bounding the distance to the feasible set We seek a bound on $\|L - L^\star\|_d$ in an appropriate norm, which we combine with the bound on $\|\mathcal{E}\|_F$ in (18) to bound the error on the objective function. Since the bound on $\mathcal{E}$ is in the Frobenius norm, ideally we set $\|\cdot\|_d = \|\cdot\|_F$, but bounds in any norm will work using $\|L\|_F \leq \sqrt{\text{rank}(L)}\|L\|$ and $\|L\|_F \leq \|L\|_*$.

Letting $F(L) = \|L\|_* + \varphi(L)$, we first bound $F(L^\star)$ by computing $F(L)$ for explicit choices of L. To be concrete, in this section we assume $r(S) = \|S\|_1$ (and again choose $\lambda_S = 1$ for simplicity) and a squared-quadratic loss function. Choosing L such that $\mathcal{A}(L) = b$ means that $S = 0$ in the definition of φ (5), and hence $F(L) = \|L\|_*$. Another choice is $L = 0$, and then choosing $S = 0$ (which may be sub-optimal) leads to $F(L) \leq \frac{1}{2}\|b\|^2$. A third choice is to explicitly compute $F(L)$ at the iterates in the algorithm and record the best.

Denoting $F^\star = \min_L F(L)$ and using $F^\star \leq F_{\text{bound}}$, non-negativity of φ immediately implies $\|L^\star\|_* \leq F_{\text{bound}}$. Hence $\|L - L^\star\|_F \leq \|L\|_F + \|L^\star\|_F \leq \|L\|_F + \|L^\star\|_* \leq \|L\|_F + F_{\text{bound}}$, and $\|L\|_F$ is explicitly computable.

[1] The orthogonal matrices U_1 and V_1 should not be confused with the variables U and V that are used as a factorization of L, as U and V are not necessarily orthogonal. In fact, U_1 and V_1 can be efficiently computed from the factorization $L = UV^T$ by taking the QR-decompositions $U = Q_UR_U$ and $V = Q_VR_V$ so $L = Q_U(R_UR_V^T)Q_V^T$. Perform a SVD on the small inner matrix to write $(R_UR_V^T) = U_R\Sigma V_R^T$, and hence Σ are the nonzero singular values of L and $U_1 = Q_UU_R$ and $V_1 = Q_VV_R$ are the corresponding left and right singular vectors.

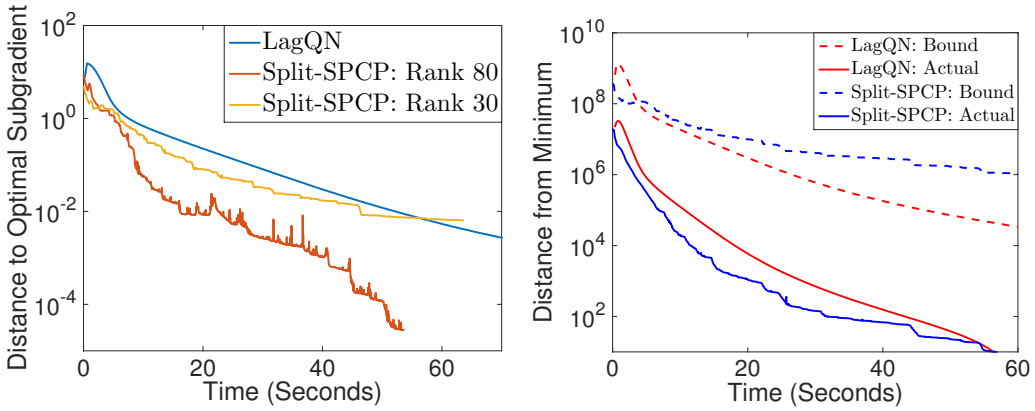

Figure 1: (Left): Distance from the optimal subgradient for the convex solver LagQN and the non-convex Split-SPCP with different rank bounds. The rank of the optimal L is 58. (Right): Bounded distance from minimal objective value, as given by $\|\mathcal{E}\|_F(\|L\|_F + F_{\text{bound}})$, with F_{bound} given by the objective value at the current iterate. The actual distance from the minimal objective value is shown for comparison. The rank bound for the Split-SPCP test is 80.

Results Figure 1 shows how the distance to the optimal subgradient decreases over time for the convex solver LagQN, the non-convex Split-SPCP with a rank bound that is too strict, and Split-SPCP with a rank bound large enough to reach the global optimum. When the rank bound is too restrictive, Split-SPCP cannot recover the optimal L. Measuring the distance to the optimal L reveals this. Figure 1 shows that the distance plateaus far from zero when the rank of L is bounded above by 30, and the rank of the optimizing L is 58. In practice, this measure can be used to indicate that the rank bound should be increased.

4 Initialization

The non-convex formulation (8) is sensitive to the point of initialization, and requires a bound on the rank of L. The performance of this solver can vary greatly for different initializations of these parameters. In this section, we present some heuristics for choosing a suitable initialization.

4.1 Dynamically Increasing k

Figure 2 demonstrates the sensitivity of the Split-SPCP program to the factor rank k. These tests were performed on the surveillance video data described in section 5, where the true rank of the low rank component is 58. We see that when $k \leq \text{rank}(L)$, Split-SPCP does not converge to the correct solution, but if k is much larger than the rank of the minimizing L, then the computation is slowed significantly.

Because our solver is oblivious to the individual dimensions of U and V, columns can be added to both matrices on the fly. Dynamically updating the rank bound can help when (8) is initialized with a k that is too small. Suppose a solver reaches some convergence criterion at iteration i. To see if this is a solution to (6), we add a single column to U and V:

$$\begin{bmatrix} U_i & u \end{bmatrix} \begin{bmatrix} V_i & v \end{bmatrix}^T = L_i + uv^T, \tag{23}$$

and observe whether this rank-one allows for a lower objective value or certificate value. Dynamically increasing k, even aggressively, is still more efficient than overestimating k from the start.

4.2 Choosing U_0 and V_0

Choosing an initialization point (U_0, V_0) has received considerable attention in factorized low rank recovery problems. Most of this work has centered around matrix completion using alternating minimization, see, for example, Hardt (2013); Keshavan & Montanari (2010); Keshavan et al. (2009, 2010); Jain et al. (2013); Udell

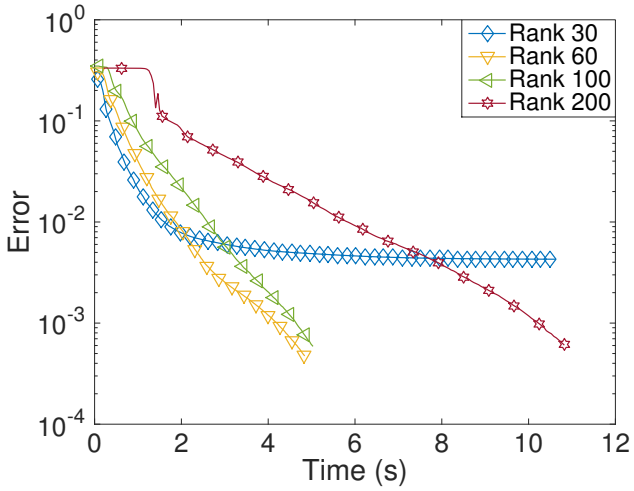

Figure 2: Varying the rank bound k in Split-SPCP affects computation time and the accuracy of the final solution. One marker represents 10 iterations.

et al. (2015). Provably good choices for U_0 and V_0 are $U_0 = U\Sigma^{\frac{1}{2}}$ and $V_0 = V\Sigma^{\frac{1}{2}}$, where U and V come from the SVD of $\mathcal{A}^\dagger X$, where X is the data-matrix, $\mathcal{A}$ is a linear operator, and $\dagger$ denotes the Moore-Penrose pseudo-inverse. Hardt (2013) showed that these initial points lie in the basin of attraction of globally optimal minima, yielding fast convergence to an optimal solution.

As mentioned in Udell et al. (2015), it is sometimes not necessary to perform a full SVD to form U_0, V_0. Once k is chosen, only a partial SVD is necessary to calculate the first k singular values and vectors, which can be done efficiently using the randomized SVD (rSVD) (Halko et al., 2011). Although using the rSVD is significantly faster than a full SVD, the values and vectors it returns are not always accurate, especially when singular values of the underlying data do not decay rapidly.

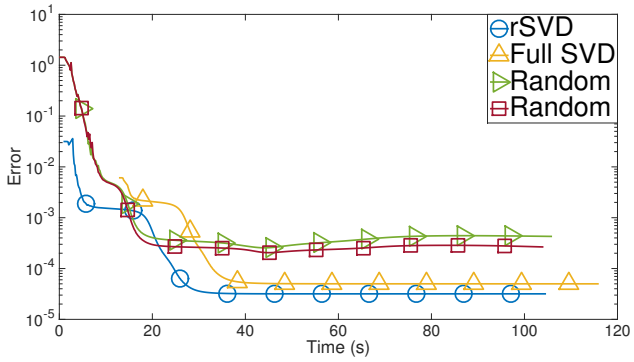

Figure 3: The performance of Split-SPCP with various initialization points. All tests were run on a 5,000 $\times$ 5,000 matrix $X = L + S$, where $\text{rank}(L) = 1000$ and $\text{sparsity}(S) = 9.43\%$. The rank bound was $k = 1{,}050$. One marker is 50 iterations.

Figure 3 shows the performance of Split-SPCP using various initial points. The tests were performed on a 5,000 $\times$ 5,000 matrix $X = L + S$, where $\text{rank}(L) = 1000$ and $\text{sparsity}(S) = 9.43\%$. The rank bound was $k = 1{,}050$. All 5,000 singular value and vector triples of X were calculated for the SVD test, and U_0 and V_0 were formed from the first $k = 1{,}050$ triples. For the rSVD test, only the first 1,050 triples were approximated, U_0 and V_0 were formed from these.

Figure 3 shows that initializing U_0 and V_0 with the first 1,050 singular value and vector triples returned by the full SVD yields the smallest initial error, but the extra computational cost is not worth the better start. The rSVD allows quicker convergence to a similar solution. The two random initializations converge to a stationary point that is not globally optimal.

5 Numerical Experiments

In this section, we present numerical experiments that illustrate our SPCP accelerations. Recall that the Split-SPCP program is given by

$$\min_{U,V} \quad \frac{\lambda_L}{2}(\|U\|_F^2 + \|V\|_F^2) + \varphi(UV^T), \tag{24}$$

where

$$\varphi(UV^T) = \min_S \frac{1}{2}\|UV^T + S - X\|_F^2 + \lambda_S\|S\|_1. \tag{25}$$

Since the objective of Split-SPCP is smooth, any first-order method can be used to solve it. Also, since the communication-heavy SVD step is no longer a limitation, we are motivated to choose the solver that is most suited for the GPU. To find this solver, we compared the performance of several first-order methods as they solved a logistic regression problem on the GPU. The time it took for different solvers to minimize a logistic loss with a Tikhonov regularization term for 10^4 labels is shown in Table 1 below. The computation time is defined as the time it took the solver to come within 10^{-8} of the optimal objective value.

Solver	Normalized Computation Time
Cyclic Steepest Descent	3.46
Barzilai-Borwein	1.24
Conjugate Gradient	1.44
Scaled Conjugate Gradient	1.53
Preconditioned Conjugate Gradient	3.20
L-BFGS (10 iterations in memory)	**1**
L-BFGS (50 iterations in memory)	1.82
BFGS	2.39
Hessian-Free Newton's Method	3.05

Table 1: First-order methods solving logistic loss program with 10^4 labels on the GPU. All times are reported as a ratio with respect to the fastest solver (L-BFGS with 10 iterations in memory). Time is recorded when the solution is within 10^{-8} of the optimal value.

From our analysis, we decided that the L-BFGS solver, storing a small number of iterations in memory, yields the best performance. The times shown in 1 depend on the stopping tolerance as well as the specific problem, but there is further evidence in the literature supporting L-BFGS as a competitively efficient solver (Wright, 1997). To solve Split-SPCP, we use L-BFGS as implemented in Schmidt (2005), with our own slight modifications so that it will run on the GPU.

The value of the objective is calculated by first computing the minimizing S in the definition of φ, which we call $S_{U,V}$, using soft thresholding. Explicitly, the soft-thresholing (or shrinkage) operator is defined so that

$$(\text{shrink}(Y, \lambda))_{i,j} = \text{sign}(Y_{i,j})(|Y_{i,j}| - \lambda|)_+,$$

where $(\cdot)_+$ is the zero operator on negative reals and is the identity operator on non-negative reals. This shrinkage operator gives us a closed-form representation of $S_{U,V}$ and $\varphi(UV^T)$:

$$S_{U,V} = \text{shrink}(X - UV^T, \lambda_S),$$
$$\varphi(UV^T) = \frac{1}{2}\|L + S_{U,V} - X\|_F^2 + \lambda_S\|S_{U,V}\|_1$$

Once $S_{U,V}$ and $\varphi(UV^T)$ are found, we can compute the objective value, $\frac{\lambda_L}{2}(\|U\|_F^2 + \|V\|_F^2) + \varphi(UV^T)$. We can also compute the gradients of the objective, which are given below. We let $f(U, V)$ denote the objective for convenience.

$$\nabla_U f = \lambda_L U + (UV^T + S_{U,V} - X)V,$$
$$\nabla_V f = \lambda_L V + (UV^T + S_{U,V} - X)^T U.$$

In Bouwmans & Zahzah (2014) and Sobral & Vacavant (2014), the authors provide a review of the fastest algorithms in the family of robust PCA and principal component pursuit for background subtraction. We chose some of the fastest of these to compare to our Split-SPCP algorithm. Each of these methods runs faster on the GPU, so we ran all of them on the GPU for our comparisons. We found that our non-convex solver is particularly well-suited for the GPU, and **outperforms all other solvers in almost every case.** This is shown on data sets from applications in section 5.1 and synthetic data in section 5.2. In section 5.3, we show that Split-SPCP on the GPU can be used for real-time background subtraction, even for high-resolution videos. Finally, in section 5.4, we demonstrate how marginalizing the sparse component in low rank models leads to an accelerated Frank-Wolfe method for optimization. We also discuss how applying marginalization to the Frank-Wolfe problem generalizes recent results published in Mu et al. (2016).

5.1 SPCP for Background Subtraction

For our background subtraction tests, we use the escalator surveillance video provided by Li et al. (2004). We want to identify the people in the video while ignoring the moving escalators and the stationary background. This problem is particularly difficult for many low rank recovery algorithms, such as PCA, because the motion of the escalator is a confounder. SPCP is less sensitive to outliers, so it can overcome this challenge. Our SPCP method and several others are compared in Figure 4. We measured performance in terms of the objective value over time. All of the algorithms were run on the GPU.

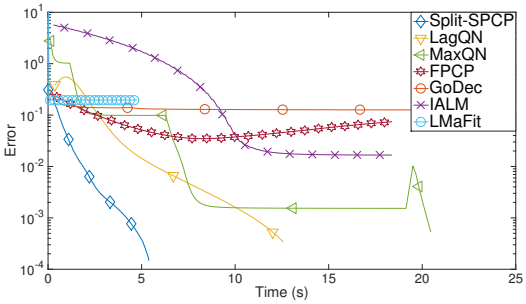

Figure 4: Comparing SPCP solvers for background subtraction. One marker corresponds to 40 iterations.

Algorithm	Reference
Split-SPCP	**This Work**
Lagrangian Quasi-Newton (LagQN)	Aravkin et al. (2014a)
Max Quasi-Newton (MaxQN)	Aravkin et al. (2014a)
Fast PCP (FPCP)	Rodrígues & Wohlberg (2013)
Go Decomposition (GoDec)	Zhou & Tao (2011)
Inexact ALM (IALM)	Lin et al. (2011)
LMaFit (LaMaFit)	Wen et al. (2010)

Table 2: References for the algorithms used in testing. Among these are the fastest models for background subtraction, as determined in Bouwmans & Zahzah (2014) and Sobral & Vacavant (2014).

To find a reference solution, we hand-tuned the parameters λ_L and λ_S in the (convex) quasi-Newton Lagrangian SPCP algorithm (LagQN) until it found a qualitatively accurate decomposition. The L we chose has rank 58, and the optimal S is 58.03% sparse. The optimal parameters for both Split-SPCP and LagQN are $\lambda_L = 115$ and $\lambda_S = 0.825$. We tuned the parameters in the other solvers to recover this solution as closely as possible. The error is measured as the normalized difference between the objective and the objective at the optimal solution. Since the solvers minimize different objectives, we calculated the Split-SPCP objective value at each iteration for every algorithm. We did not include these calculations in our time measurements. To initialize Split-SPCP, we used the first 100 singular values and vectors from the randomized SVD of the data matrix.

Several of the algorithms in Figure 4 did not converge to the same solution, despite considerable effort in parameter tuning. These algorithms were designed to quickly find approximate solutions to the SPCP problem, and might have suffered from the large amount of noise present in the data. The approximate solutions recovered by these algorithms are qualitatively different from the solution found by Split-SPCP, see Figure 5. We also found that a lower objective value generally corresponds to a qualitatively superior solution.

We measured the quality of the solutions in terms of the *Corrected Akaike Information Criterion*, or AIC_C. The AIC_C measures the fit of a statistical model to a certain data set. Given a model with p

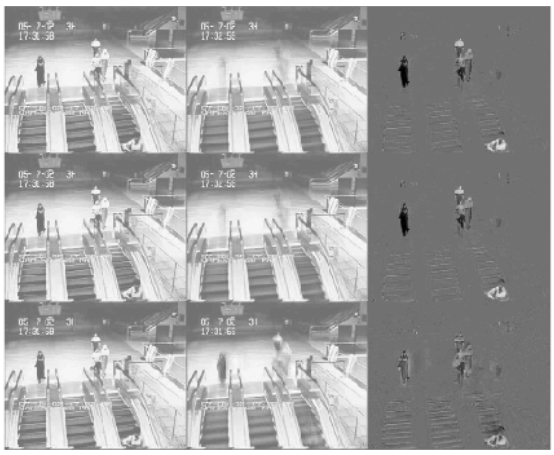

(a) X,L, and S matrices found by (from top to bottom) Split-SPCP, LagQN, and GoDec.

(b) X,L, and S matrices found by (from top to bottom) FPCP, IALM, and LMaFit.

Figure 5: Background subtraction using various SPCP solvers on surveillance video data from Li et al. (2004) (frame 10 is shown). We see that Split-SPCP, LagQN, and IALM best locate the people while ignoring the escalators.

parameters and a log-likelihood function ℓ, the value of the AIC_C is

$$\text{AIC}_\text{C} = 2(p - \log(\ell_{\max})) + \frac{2p(p+1)}{m \cdot n - p - 1},$$

where $m \cdot n$ is the size of the data set and $\ell_{\max}$ is the maximum of ℓ. The preferred statistical model is the one that minimizes the AIC_C. The AIC_C does not only favor models that maximize the likelihood function, but it also penalizes complex models with many parameters, guarding against overfitting. Using AIC_C as a measure of quality avoids both overemphasis of objective values and the inherent ambiguity of visual comparisons.

To compute the AIC_C value, we must formulate SPCP as a statistical model. It is well-known that the least-squares loss term assumes that the data is of the form $X = L + S + Z$, where the entries of Z are iid Gaussian random variables with $\mu = 0$ and variance estimated using the sample variance, $\hat{\sigma}^2 = \frac{\|X\|_F^2}{m \cdot n}$. Similarly, the ℓ_1-regularizer assumes that the entries of S are drawn iid from the Laplace distribution with mean 0 and variance estimated as $\hat{b} = \frac{\|S\|_1}{m \cdot n}$. The nuclear norm of L is the ℓ_1-norm of its singular values, so its corresponding prior assumes that the singular values of L follow the Laplace distribution. The log-likelihood function is then

$$\ell(L, S, X) = -\left(\frac{m \cdot n}{2}\right) \log(2\pi\sigma^2) - \frac{\|L + S - X\|_F^2}{2\sigma^2} - (m \cdot n)\log(2b) - \frac{\|S\|_1}{2b} - \text{rank}(L)\log(2b_*) - \frac{\|L\|_*}{2b_*}, \tag{26}$$

where σ^2, b and b_* are computed according to their respective estimator. We must also define the number of parameters of this model, which is equal to its degrees of freedom. Each term in SPCP provides the following degrees of freedom:

$$\begin{aligned} \text{rank}(L) = k \quad &\rightarrow \quad k(m + n - k) \text{ degrees of freedom,} \\ \|S\|_1 \quad &\rightarrow \quad nnz(S) \text{ degrees of freedom,} \\ \tfrac{1}{2}\|L + S - X\|_F^2 \quad &\rightarrow \quad \left(\frac{\|L+S-X\|_F^2}{\|X\|_F^2}\right)(m \cdot n) \text{ degrees of freedom.} \end{aligned}$$

Note that $nnz(S)$ counts the number of non-zero entries in S. Finding the degrees of freedom in a rank-k matrix or a sparse matrix are both standard calculations. For the loss term, we use the residual effective

degrees of freedom as an estimate of the flexibility that it introduces. This is detailed further in Loader (1999). The number of parameters p is equal to the total degrees of freedom.

The AIC_C values for the solutions shown in Figure 5 are listed in Table 3. The AIC_C value for the reference solution is listed under "oracle." Recall that the reference solution was found using the LagQN solver with a tight tolerance. All the other values in the table correspond to solutions that met comparable tolerances. Also, since the IALM model does not use an ℓ_1-norm regularizer, many of the values in the returned S matrix were small but not exactly zero. To make accurate comparisons, we applied the shrinkage operator to the S matrix returned by IALM to set small values equal to zero before calculating the degrees of freedom. The values for IALM and LMaFit are high because these solvers are not generally robust to large amounts of noise. With this metric as well, we see that Split-SPCP discovers the best solution, and in a much shorter time compared to the other algorithms.

Solver	AIC_C ($\times 10^6$)
Oracle	*7.37*
Split-SPCP	**7.64**
GoDec	10.95
LagQN	11.30
MaxQN	11.64
FPCP	15.70
IALM	2.43×10^7
LMaFit	3.99×10^9

Table 3: Degrees of freedom in solutions returned by various solvers.

Split-SPCP has an advantage because it parallelizes well on GPU architectures. This is due to the relative simplicity of the L-BFGS algorithm, see also Feia et al. (2014). The algorithm avoids complex linear algebraic decompositions that require heavy communication, such as the QR and SVD. Figure 6 shows that although Split-SPCP on the CPU is slower than LagQN on the CPU, Split-SPCP enjoys enormous speedup when it is implemented on the GPU.

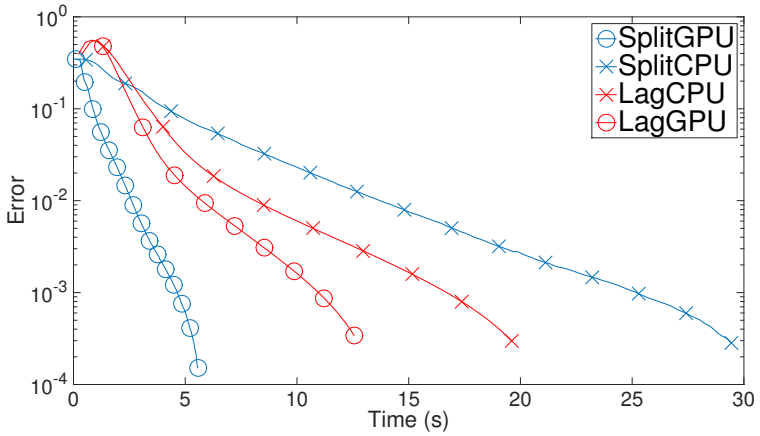

Figure 6: Split-SPCP and LagQN on the CPU and the GPU. Tests were performed on the synthetic data described in section 5.2

5.2 Synthetic Data

To test Split-SPCP further, we created a synthetic problem consisting of a noisy data-matrix that can be represented as the sum of a low rank and a sparse component. We created two $1{,}000 \times 150$ random matrices with each entry drawn from independently from the univariate normal distribution, and defined our low rank reference matrix, L_{ref}, as the product of these two matrices. To form the sparse reference matrix,

S_{ref}, we filled 50% of a $1{,}000 \times 1{,}000$ matrix with numbers drawn independently from the univariate normal distribution, and the other entries we set equal to zero. We defined X to be the sum $L_{ref} + S_{ref}$ with added white Gaussian noise, so that $\frac{\|L_{ref}+S_{ref}-X\|_F}{\|X\|_F} = 8.12 \times 10^{-5}$. Compared to data sets that are normally encountered in applications of SPCP, the sparse component of S has a large number of non-zero entries, the rank of L is large, and the noise in X is small.

We initialized Split-SPCP with the first 200 singular value and vector triples of X, so that k was larger than the rank of L. These triples were found using the rSVD. The tuning parameters were set to $\lambda_L = 2.95$ and $\lambda_S = 0.1$. As before, we measured performance based on the normalized difference from the true Split-SPCP objective value. The results are shown in Figure 7.

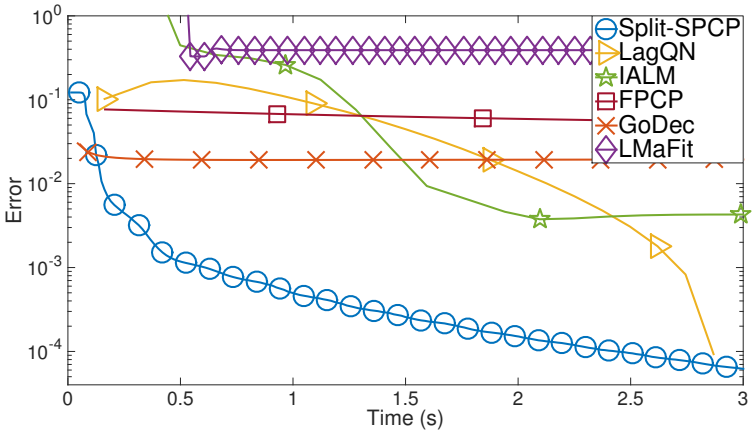

Figure 7: Performance of SPCP solvers on synthetic data. One marker represents six iterations.

We see that FPCP and GoDec discover approximate solutions quickly, but the approximations are not within 10^{-2} of the true objective. LMaFit exhibits poor performance. This could be due to the fact that S is about 50% sparse, and LMaFit struggles with problems that have relatively many non-zero entries (Wen et al., 2010). Split-SPCP significantly outperforms all solvers in speed and accuracy, except for LagQN. While Split-SPCP is much faster than LagQN initially, LagQN catches up in later iterations.

5.3 Dimensional Scaling and Real-Time Video Processing

One useful application of accelerated low rank recovery algorithms is real-time background subtraction in video. The experiments in section 5.1 show that Split-SPCP is faster and more accurate than the methods reviewed in Bouwmans & Zahzah (2014) and Sobral & Vacavant (2014) for background subtraction, and in this section, we apply Split-SPCP to the real-time processing of high-resolution videos with standard frame rates. We also explore how Split-SPCP scales as the resolution and number of frames increases.

To measure the performance of Split-SPCP on more diverse video datasets, we used videos from the Background Model Challenge (Vacavant et al., 2012). The experiments shown in Figure 8 used a video with 240×320 pixels per frame and 32,964 frames, which is substantially larger than the low-resolution escalator video with 130×160 pixels per frame and only 200 frames. We unfolded the three-dimensional video data into a matrix with each column containing a frame, and we partitioned the matrix columnwise into blocks comprising a fixed number of frames. With the rank bound fixed at 60, we then recorded the computation time required to perform 200 iterations of Split-SPCP on one block as a function of the block size. Larger block sizes yield better solutions, but they also require more time for SPCP to converge. We found that after 200 iterations, Split-SPCP had converged to within 10^{-3} of the optimal objective value for all block sizes, and the returned solution was qualitatively no better than the optimal solution.

Figure 8 shows how Split-SPCP scales as the video resolution and the number of frames per block increase. These two qualities correspond to the number of rows and the number of columns of each block. We increased the resolution by rescaling each frame so that the aspect ratio stayed close to constant, and we used interpolation to impute the missing pixel values. It is easy to discern the $\mathcal{O}(n)$ scaling as the number of frames (columns) is increased. Increasing the resolution affected the computation time linearly, but beyond a

"critical resolution" of 332×446 pixels per frame, the slope of the linear scaling underwent a shift in regime. This linear scaling is expected because the bottleneck of Split-SPCP is the matrix-matrix multiply UV^T, which requires $\mathcal{O}(mnk)$ operations, so with k fixed, the computation time of Split-SPCP grows linearly with the number of data points. The results of the background subtraction are shown in Figure 9.

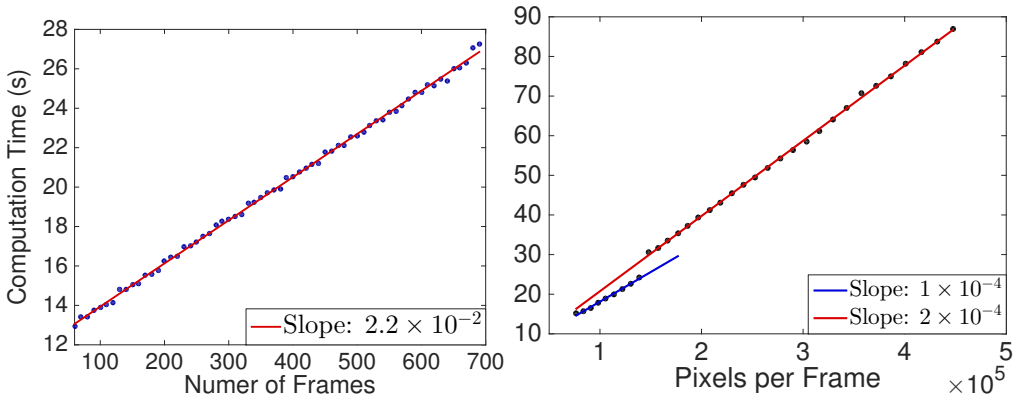

Figure 8: The scaling of Split-SPCP (left) as the number of columns (or frames) increases and (right) as the number of rows (or pixels per frame) increases. Each point is the average of twenty trials. For the column-scaling tests, the resolution was fixed at 240×320 pixels per frame, and for the row-scaling tests, the number of frames in a block was fixed at 100. The tuning parameters were set to $\lambda_L = 200, \lambda_S = 5$.

These experiments suggest conditions under which Split-SPCP can decompose a video in real-time. Assuming a resolution of 240×320 pixels per frame and a frame rate of 24 frames-per-second, which is standard in many applications of video processing, the video must be partitioned blocks of about 600 frames or more. With this partitioning, the algorithm will have finished decomposing the given block before the next block is recorded. This number increases linearly with resolution, following the trend line given in the rightmost plot of Figure 8. Following these guidelines, Split-SPCP can decompose the videos from Vacavant et al. (2012) in real-time without sacrificing quality. An example is shown in Figure 9.

5.4 Accelerating Frank-Wolfe

Our approach to accelerating low rank recovery models is closely connected to an approach presented in Mu et al. (2016) to accelerate the Frank-Wolfe method for RPCA. The Frank-Wolfe method, first proposed in Frank & Wolfe (1956), has recently gained attention for its ability to yield scalable algorithms for optimization problems involving structure-encouraging regularizers (Jaggi, 2013; Mu et al., 2016). Frank-Wolfe is a method of optimizing a Lipschitz-differentiable function over a convex set by solving a sequence of linear subproblems. For example, if f is Lipschitz-differentiable and $\mathcal{C}$ is a convex set, Frank-Wolfe solves the problem

$$\min_{x \in \mathcal{C}} \quad f(x) \tag{27}$$

by linearizing f about each iteration x_k: $f(y) \approx f(x_k) + \langle \nabla f(x), y - x \rangle$, minimizing the linear problem to find a decent direction $y_k - x_k$, and then stepping in this direction with a step size η: $x_{k+1} = x_k + \eta(y_k - x_k)$. Traditionally, $\eta = \frac{2}{k+2}$ is fixed for convergence guarantees, but better performance can be achieved with adjustable step sizes Mu et al. (2016). A summary of this process is shown in Algorithm 1.

The Frank-Wolfe method has recently become popular in machine learning, statistics, image processing, and related fields due to the simple form of the linear subproblem when handling structure-inducing regularizers (Jaggi, 2013; Jaggi & Sulovsk, 2010; Harchaoui et al., 2015). For example, for any optimization problem involving the nuclear norm, we can write an equivalent problem where the nuclear norm term is represented in the constraint (Jaggi & Sulovsk, 2010; Mu et al., 2016):

$$\min_X \quad f(X) + \lambda \|X\|_* \Leftrightarrow \min_{\|X\|_* \leq t} \quad f(X) + \lambda_L t. \tag{28}$$

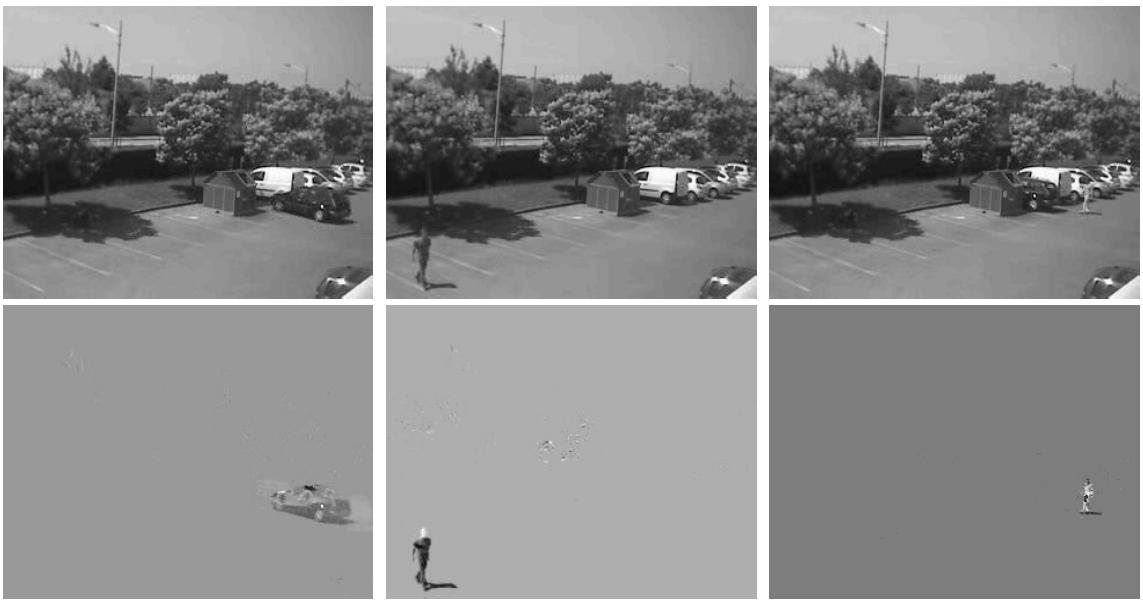

Figure 9: Real-time background subtraction using Split-SPCP. Split-SPCP correctly identifies (left) a car in the middle-ground, (center) a person in the foreground, and (right) a person in the background. The resolution is 240×320 pixels per frame, 600 frames were decomposed at one time, and only 200 iterations were allowed.

Algorithm 1: Frank-Wolfe Method to Solve (27) (e.g. Frank & Wolfe (1956); Jaggi (2013); Mu et al. (2016))

Input : $x_0 \in \mathcal{C}$

while (Not Converged) **do**

- $y_k \leftarrow \arg\min_{y \in \mathcal{C}} \langle y, \nabla f(x_k) \rangle$
- $\eta = \frac{2}{k+2}$
- $x_{k+1} \leftarrow x_k + \eta(y_k - x_k)$

When $\mathcal{C}$ in Algorithm 1 is the scaled nuclear norm unit ball, then the linear subproblem corresponds to finding the leading eigenspace of the iterate L_k. This is much cheaper than computing the full SVD at each iteration, which is required by proximal methods.

However, it is often difficult to scale the Frank-Wolfe method for problems involving ℓ_1 regularization. For these problems, each Frank-Wolfe iteration updates iterate S_k with only a single non-zero element (Mu et al., 2016). This makes Frank-Wolfe unfeasible when the sparse component of the data-matrix has several thousand non-zero elements or more. The authors of Mu et al. (2016) adapt the Frank-Wolfe method to solve SPCP, using the traditional Frank-Wolfe update for the low rank component and using a projected gradient update for the sparse component. Their adaptations allow for significant speedup and better convergence compared to the traditional Frank-Wolfe.

The techniques used in Mu et al. (2016) to accelerate the update of the sparse component can be seen as a special case of the marginalization we present in this work. With the sparsely structured component marginalized, the program depends only upon the low rank component, so the benefits of the Frank-Wolfe scheme are preserved and its drawbacks are negated. Algorithm 2 shows how marginalization can be used to extend the adaptations presented in Algorithms 5 and 6 of Mu et al. (2016). Algorithm 2 solves the following program, which is equivalent to SPCP:

$$\min_{\|L\|_* \leq t} \quad \lambda_L t + \min_S \frac{1}{2} \|\mathcal{P}_\Omega(L + S - X)\|_F^2 + \lambda_S \|S\|_1. \tag{29}$$

This problem is smooth, and as mentioned earlier, the constraint set is amenable to acceleration by applying the linear subproblems introduced in the Frank-Wolfe scheme. Each iteration of Algorithm 2, then, updates the low rank component and the nuclear norm bound, t, while the sparse component remains only implicitly defined. More details on this approach can be found in Mu et al. (2016).

While Algorithm 2 is similar to the methodology in Mu et al. (2016), fully marginalizing the sparse component eliminates the linear subproblem required in Mu et al. (2016) to update the iterate S_k. We show the performance of this adapted Frank-Wolfe scheme on SPCP in Figure 10. Although its performance is better than the traditional Frank-Wolfe (see Mu et al. (2016)), it is still much slower than quasi-Newton methods applied to the smoothed problem.

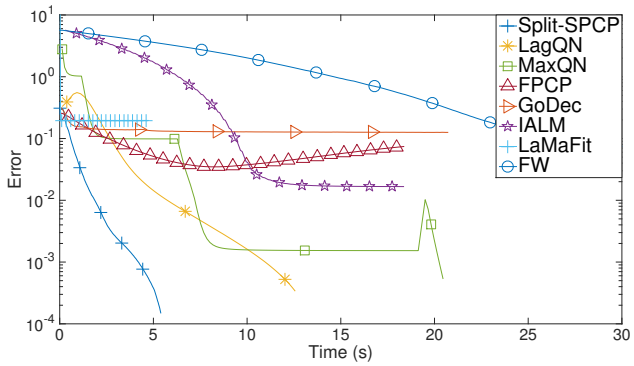

Figure 10: Comparing Frank-Wolfe with marginalization (denoted "FW") with other SPCP solvers for background subtraction on the escalator video of Li et al. (2004). One marker corresponds to 40 iterations.

Algorithm 2: Frank-Wolfe Method with Marginalization

Input : $\mathbf{L_0}, \lambda_L, \lambda_S, t_0 = 0, U_0 = \frac{1}{2\lambda_L}\|\mathcal{P}_\Omega(\mathbf{L_0} + \mathbf{S_0} - \mathbf{X})\|_F^2$

while Not Converged **do**

- $\mathbf{S_k} \leftarrow \text{shrink}(\mathbf{X} - \mathbf{L_k}, \lambda_S)$
- $\mathbf{Y_k} \leftarrow \underset{\|\mathbf{Y}\|_* \leq 1}{\arg\min} \left\langle \mathcal{P}_\Omega(\mathbf{L_k} + \mathbf{S_k} - \mathbf{X}), \mathbf{Y} \right\rangle$
- **if** $\lambda_L \geq -\left\langle \mathcal{P}_\Omega(\mathbf{L_k} + \mathbf{S_k} - \mathbf{X}), \mathbf{Y_k} \right\rangle$ **then**
  - $\mathbf{V_k} \leftarrow \mathbf{0}$
  - $V_{t_k} \leftarrow 0$
- **else**
  - $\mathbf{V_k} \leftarrow U_k \mathbf{Y_k}$
  - $V_{t_k} \leftarrow U_k$
- $\eta^\star \leftarrow \underset{\eta \in (0,1)}{\arg\min} \quad \frac{1}{2}\|\mathcal{P}_\Omega((1-\eta)\mathbf{L_k} + \eta\mathbf{V_k} + \mathbf{S_k} - \mathbf{X})\|_F^2 + (1-\eta)\lambda_L t_k + \eta\lambda_L t_k$
- $\mathbf{L_{k+1}} \leftarrow (1-\eta^\star)\mathbf{L_k} + \eta^\star \mathbf{V_k}$
- $t_{k+1} \leftarrow (1-\eta^\star)t_k + \eta^\star V_{t_k}$
- $U_{k+1} \leftarrow \frac{1}{2\lambda_L}\|\mathcal{P}_\Omega(\mathbf{L_0} + \mathbf{S_0} - \mathbf{X})\|_F^2 + t_k$

6 Low Rank Recovery Models with a Parallel rSVD

In section 5, we saw that solving low rank recovery models without the SVD step required by proximal methods allows the GPU to offer greater acceleration. In this section, we explore whether a communication-avoiding rSVD can provide the convex problem (3) with similar acceleration on the GPU.

Many problems in data analysis involve matrices of a "tall-skinny" structure. Tall-skinny matrices arising in applications could have twice as many rows as columns, or several-thousand-times more rows than columns

(Demmel et al., 2012). In the psychometric data that we consider in section 6.4, the matrices have about 230-times more rows than columns. The surveillance video from the previous section has dimensions $20{,}800 \times 200$, so it is also tall-skinny. The prevalence of this structure is due to the fact that high-dimensional data sets must be "unwrapped" into a two-dimensional matrix before matrix-based algorithms can analyze the data. Data from video are three-dimensional: there are two dimensions of space, and one dimension of time. When applying a low rank recovery model to this data, each frame of the video is stored as one column of the matrix, so a two-hundred frame video with a resolution of 120×160 pixels per frame produces a $20{,}800 \times 200$ matrix for decomposition. The problem is exacerbated when data is even higher dimensional, as it is for fMRI data.

Unfortunately, tall-skinny matrices pose a challenge for matrix decompositions, especially on computational architectures where communication is costly (such as a GPU) (Demmel et al., 2012). This is why transferring low rank recovery models to the GPU is a challenge: to minimize the nuclear norm, an SVD is computed at every step, and this is extremely costly on the GPU.

Even for problems involving matrices that are not originally tall-skinny, many popular algorithms for data-analysis produce tall-skinny matrices during intermediate calculations, which significantly slows computation. The randomized SVD (Halko et al., 2011) is one of these algorithms. The rSVD creates a tall-skinny matrix structure that amplifies communication costs, even if the input matrix is square. The effect of this increased communication is notable: on the CPU, the rSVD step in LagQN accounts for about 55% of the algorithm's total running time, but when this solver is run on the GPU, the rSVD accounts for about 90% of SPCP's total running time. The following sections demonstrate how the Tall-Skinny QR algorithm of Demmel et al. (2012) can provide a convex option for accelerated low rank recovery models. While TSQR has been used to accelerate the SVD step in RPCA for tall-skinny matrices Anderson et al. (2010), we extend this idea to the rSVD algorithm on general matrices.

6.1 Motivation

A basic description of the rSVD algorithm is presented in Algorithm 3. To decompose the $m \times n$ matrix

Algorithm 3: Randomized Singular Value Decomposition (rSVD) (Halko et al., 2011)

Input : Matrix X, rank estimate k, and power-parameter q
Output: U, Σ, V so that $X = U\Sigma V^T$

1 Form Ω, an $n \times k$ Gaussian matrix
2 $Y_0 \leftarrow X\Omega$, so that Y is an $m \times k$ tall-skinny matrix
3 $Q_0 \leftarrow \texttt{qr}(Y_0)$
4 **for** $j \in \{1, 2, \cdots q\}$ **do**
5 $\quad \tilde{Y}_j \leftarrow X^T Q_{j-1}$
6 $\quad \tilde{Q}_j \leftarrow \texttt{qr}(\tilde{Y}_j)$
7 $\quad Y_j \leftarrow X\tilde{Q}_j$
8 $\quad Q_j \leftarrow \texttt{qr}(Y_j)$
9 $Q \leftarrow Q_q$
10 $B \leftarrow Q^T X$, (where B is a small $k \times n$ matrix)
11 $[\hat{U}, \Sigma, V] \leftarrow \texttt{svd}(B)$
12 $U \leftarrow Q\hat{U}$

X, the rSVD algorithm first multiplies X on the right by the $n \times k$ random matrix Ω to project X onto a low-dimensional subspace. This projection creates an $m \times k$ matrix Y, where $k \ll m$, The orthogonalization steps (3 through 8) perform QR decompositions on these tall-skinny matrices. If the singular values of X decay slowly (which is often the case in applications), then several power iterations are performed, which increases the number of QR decompositions. Regardless of input X, Y is always tall-skinny when the target rank k is small.

The tall-skinny structure notably slows the QR decomposition on the CPU, and when performing this algorithm on the GPU, the sub-optimality of the tall-skinny structure is amplified (Demmel et al., 2012).

To avoid this problem, we replace the QR decompositions in the rSVD algorithm with a tall-skinny QR decomposition developed in Demmel et al. (2012).

6.2 TSQR Algorithm

It is important to see why tall-skinny matrix structures are so hazardous in the first place. This analysis can be found in Demmel et al. (2012), but we include it here for completeness. Consider the Householder QR decomposition on the tall-skinny matrix in Figure 11:

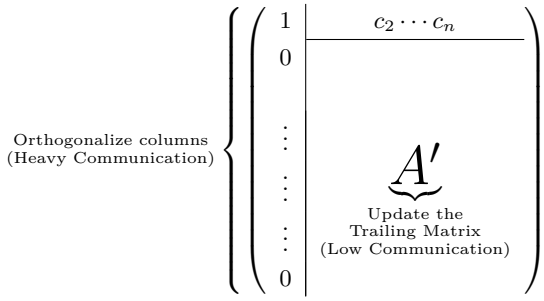

Figure 11: The first iteration in the Householder QR Decomposition of a Tall-Skinny Matrix.

Orthogonalization of columns is a particularly communication-heavy task, especially when compared to the matrix-update step that can be performed using BLAS3 matrix-matrix multiplies that are easily parallelizable. As the columns of the matrix get taller, the Householder QR must spend more time orthogonalizing columns and less time performing the matrix-update. This causes deceleration in CPUs, and in GPUs, where communication between cores is even more time consuming, this causes significant inefficiency. The TSQR decomposition reduces this inefficiency. An outline of the algorithm is as below.

TSQR Algorithm (Demmel et al., 2012):

1. Divide the matrix vertically into 2^n blocks, which we illustrate with $n = 2$:

$$A = \begin{bmatrix} A_{1,1} \\ A_{1,2} \\ A_{1,3} \\ A_{1,4} \end{bmatrix}$$

2. Perform a QR decomposition on each block in parallel:

$$\begin{bmatrix} A_{1,1} \\ A_{1,2} \\ A_{1,3} \\ A_{1,4} \end{bmatrix} = \left[\begin{array}{c|c|c|c} Q_{1,1} & & & \\ \hline & Q_{1,2} & & \\ \hline & & Q_{1,3} & \\ \hline & & & Q_{1,4} \end{array}\right] \left[\begin{array}{c} R_{1,1} \\ \hline R_{1,2} \\ \hline R_{1,3} \\ \hline R_{1,4} \end{array}\right]$$

3. Concatenate the matrix of R's vertically, forming the matrices $A_{2,i}$, for $i = 1, 2, 3, \cdots, 2^{n-1}$:

$$\begin{bmatrix} A_{2,1} \\ A_{2,2} \end{bmatrix} \leftarrow \left[\begin{array}{c} \begin{pmatrix} R_{1,1} \\ R_{1,2} \end{pmatrix} \\ \hline \begin{pmatrix} R_{1,3} \\ R_{1,4} \end{pmatrix} \end{array}\right]$$

4. Repeat steps 2 through 4 for the matrices $A_{2,i}$. Do this process $n - 1$ times. For $n = 2$, we form the matrices $Q_{2,1}, Q_{2,2}, R_{2,1}$, and $R_{2,2}$:

$$\begin{bmatrix} A_{1,1} \\ A_{1,2} \\ A_{1,3} \\ A_{1,4} \end{bmatrix} = \left[\begin{array}{c|c|c|c} Q_{1,1} & & & \\ \hline & Q_{1,2} & & \\ \hline & & Q_{1,3} & \\ \hline & & & Q_{1,4} \end{array}\right] \left[\begin{array}{c|c} Q_{2,1} & \\ \hline & Q_{2,2} \end{array}\right] \left[\begin{array}{c} R_{2,1} \\ \hline R_{2,2} \end{array}\right]$$

5. Concatenate the matrix of R's to form $A_{3,1}$:

$$A_{3,1} \leftarrow \begin{bmatrix} \begin{pmatrix} R_{2,1} \\ R_{2,2} \end{pmatrix} \end{bmatrix}$$

6. Performing the final QR decomposition on $A_{3,1}$ we have

$$\begin{bmatrix} A_{1,1} \\ A_{1,2} \\ A_{1,3} \\ A_{1,4} \end{bmatrix} = \left[\begin{array}{c|c|c|c} Q_{1,1} & & & \\ \hline & Q_{1,2} & & \\ \hline & & Q_{1,3} & \\ \hline & & & Q_{1,4} \end{array}\right] \left[\begin{array}{c|c} Q_{2,1} & \\ \hline & Q_{2,2} \end{array}\right] Q_{3,1} R_{3,1}$$

7. The upper-triangular matrix R is equal to $R_{3,1}$, and the orthogonal matrix Q is the product

$$Q = \left[\begin{array}{c|c|c|c} Q_{1,1} & & & \\ \hline & Q_{1,2} & & \\ \hline & & Q_{1,3} & \\ \hline & & & Q_{1,4} \end{array}\right] \left[\begin{array}{c|c} Q_{2,1} & \\ \hline & Q_{2,2} \end{array}\right] Q_{3,1},$$

but since R is upper-triangular, it is often more efficient to calculate $Q = AR^{-1}$.

The QR decomposition in steps 2,4, and 6 is generally chosen to be Householder's, but any method could be used without changing the high-level flow of the algorithm. Figure 12 clearly summarizes the binary-tree structure of TSQR. The optimal number of vertical partitions depends on the height-to-width ratio of the input matrix and on the particular architecture that is used to execute the algorithm.

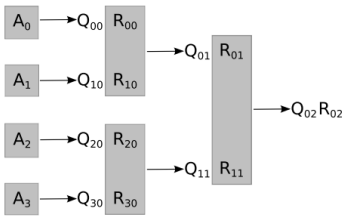

Figure 12: The TSRQ algorithm's binary tree structure. Figure from Demmel et al. (2012).

It is easy to see the speedup in the early steps of the TSQR algorithm due to the vertical partitioning and parallelization, but it is not obvious that the final steps of TSQR maintain the algorithm's efficiency. In the first step, the vertical partitions ameliorate the tall-skinny communication problem, and the "short" QR decompositions can be performed in parallel, which is another source of speedup. However, as we descend the TSQR tree, the matrices become taller and the parallelization disappears, so these steps become more costly. As long as the input matrix has sufficiently more rows than columns, then the intermediate R submatrices shrink, so that the final QR performed without parallelization is smaller than the original matrix. This means that the final QR is still inexpensive compared to performing a traditional QR on the original matrix.

To exhibit this shrinking, suppose we perform a two-split TSQR on a 64×4 matrix. The first step is to perform QR's on the two 32×4 matrices that result from the split. This produces two 4×4 matrices, $R_{1,1}$ and $R_{1,2}$. The final step is to concatenate these two matrices and perform a QR decomposition on the resulting 8×4 matrix. In total, the number of operations required to do this TSQR decomposition is equal

to the number of operations required to find the QR of one 32×4 matrix and one 8×4 matrix, and this costs less than performing a QR on the original 64×4 matrix. In general, if the input matrix is size $m \times n$ and the number of splits in the TSQR algorithm is s, then the final QR will be less expensive than the initial QR if $m \geq n^s$.

6.3 Numerical Experiments on rSVD

To test the ability of TSQR to accelerate rSVD, we built implementations of TSQR that run on (1) multiple CPU cores and on (2) multiple blocks of a single GPU card. The multiple CPU code uses MATLAB's Parallel Computing Toolbox, while the GPU code uses Magma libraries (Tomov et al., 2010). Our code can also be used to run on multiple GPU cards, but due to limited resources, we exclude this case from our tests.

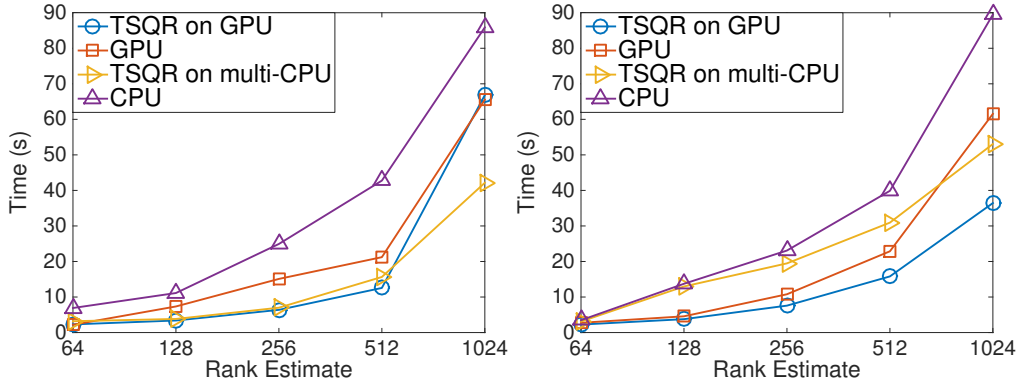

Figure 13: Timing the rSVD algorithm for various rank estimates and matrix sizes: on the left, the matrix is $65{,}536 \times 8{,}192$. On the right, the matrix is $131{,}072 \times 4{,}096$ (twice the number of rows and half the number of columns).

We test four different rSVD implementations: rSVD on the CPU, rSVD on the GPU, rSVD on the GPU with a multi-CPU-core TSQR, and rSVD on the GPU with a single-GPU TSQR. The matrices in our tests were randomly generated with each entry drawn independently from the univariate normal distribution, so they were all of full rank. The rank estimate is the number of columns that the matrix Y has at the time of the QR decomposition, so the speedup that TSQR offers would be the same even if the test matrices were square. Both tests (and both the multi-CPU and GPU TSQR algorithms) employed sixteen vertical partitions for the TSQR step, which we determined to be near optimal through testing. We see from the tests that both TSQR on the GPU and on multi-CPU architectures provide significant speedup. TSQR on the GPU is most effective for extremely tall-skinny matrices, while TSQR on the CPU outperforms its GPU counterpart when the matrix is closer to square.

6.4 Numerical Experiments on fMRI Brain Scan Data

Functional MRI brain scan data sets are inherently four-dimensional, so unwrapping the dataset into a two-dimensional array creates a tall-skinny matrix with hundreds of times more rows than columns. The data set used in the following experiments has 230-times more rows than columns. Also, the enormity of these data sets makes low rank recovery models intractable without parallelization. Figure 14 shows the total time spent performing SPCP on one of these data sets, as well as the total time spent performing rSVDs, performing the QR step of the rSVD, and moving data (labeled as the "overhead" cost). The results of this test are included in Figures 15 and 16.

The datasets used for testing were taken from a study of the human brain's response to uncomfortable stimuli, specifically, the feeling of submerging a hand in hot or cold water. Analyzing these scans to find common neurological responses to these stimuli is difficult due to the enormous amount of error in the data. There is uniformly distributed noise due to constant ancillary physiological activity, and there are also sparsely distributed groups representing neurological structures that should all exhibit the same behavior.

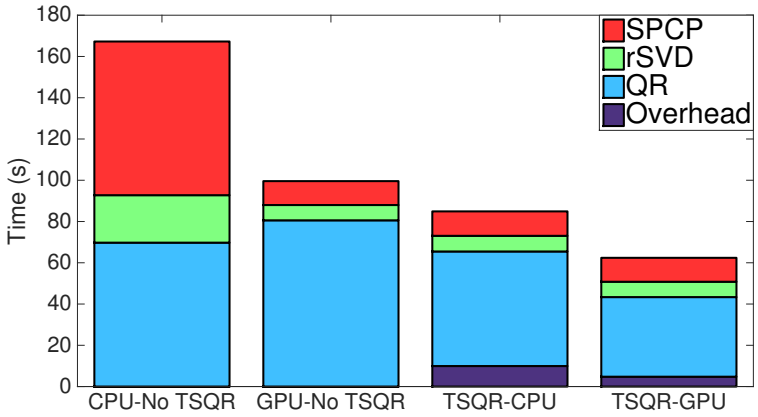

Figure 14: The total time spent performing various tasks during a run of SPCP on fMRI brain scan data. The test "TSQR-CPU" uses a hybridized architecture, running TSQR on multiple CPU's but running the rest of SPCP on the GPU. Both tests "GPU-No TSQR" and "TSQR-GPU" were run entirely on the GPU.

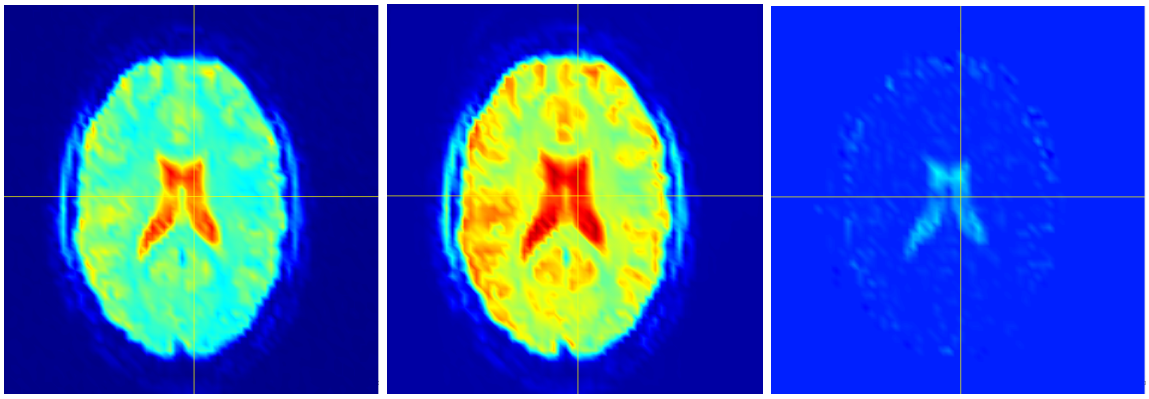

Figure 15: The results of SPCP on fMRI brain scan data. Activity is measured in Blood-oxygen-level dependent (BOLD) signal. (From left to right: original image, low rank image recovered by SPCP, and sparse image recovered by SPCP.)

The ventricles, for example, are filled with cerebrospinal fluid (CSF), which does not contribute to neurological communication, so they should not be active. All signals observed in the ventricles should be treated as sparsely structured outliers. SPCP removes the uniform noise and, most remarkably, correctly identifies signals in the brain's ventricles as outliers. In Figure 15, the largest ventricles are the two structures in the center of the brain. The rightmost image shows that the majority of the noise contained in S is from these ventricles.

The other two major components of the brain are white and gray matter. The activity we are hoping to observe takes place in the gray matter, so ideally, SPCP would remove most signals from the white matter regions. However, the regions of white matter are more difficult to distinguish than the regions of CSF, and SPCP removes about equal amounts of noise from the white matter as it does from the gray. If we let S_{gm} be the gray-matter component of S, and define S_{wm} and S_{csf} similarly, Figure 16 shows the average BOLD signal in S_{gm}, S_{wm}, and S_{csf} for each frame in time. These data were normalized by the average original signal in the corresponding regions.

It is clear that S_{csf} contains more signal than the other two regions. For S_{gm}, SPCP detects noise in only the first 100 time slices. The removed signal from the white matter is more distributed over time, and the total amount of noise in S_{gm} and S_{wm} is comparable. These results suggest that SPCP correctly identifies outliers in the fMRI data, especially within the regions of CSF.

Because of the size and tall-skinny structure of fMRI brain scan data, this analysis using SPCP is only feasible using parallel programming. Figure 14 shows how incorporating the TSQR algorithm into SPCP on

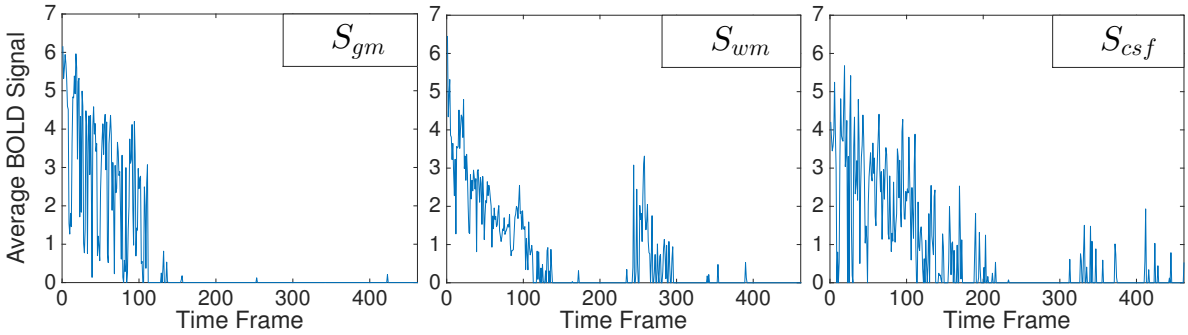

Figure 16: The average BOLD signal in different regions of S. The averages were normalized by the average original signal in the corresponding regions.

the GPU more than doubles the speed of the algorithm. While the GPU offers enormous speedup over the CPU implementation of SPCP, the rSVD step on the GPU is about 90% of the algorithm by time, and it is only about 55% of the algorithm on the CPU. This discrepancy is ameliorated by TSQR.

6.5 Comparing the Non-Convex Model to the Communication-Avoiding Approach

Although reducing communication in the low rank recovery algorithm provides noticeable speedup, the non-convex approach is often even faster. In Figure 17, we compare our non-convex SPCP algorithm to a convex solver using TSQR for the rSVD step. Both solvers were decomposing the same 106,496 $\times$ 462 brain-scan data-matrix used in section 6.4, and both were run on the GPU.

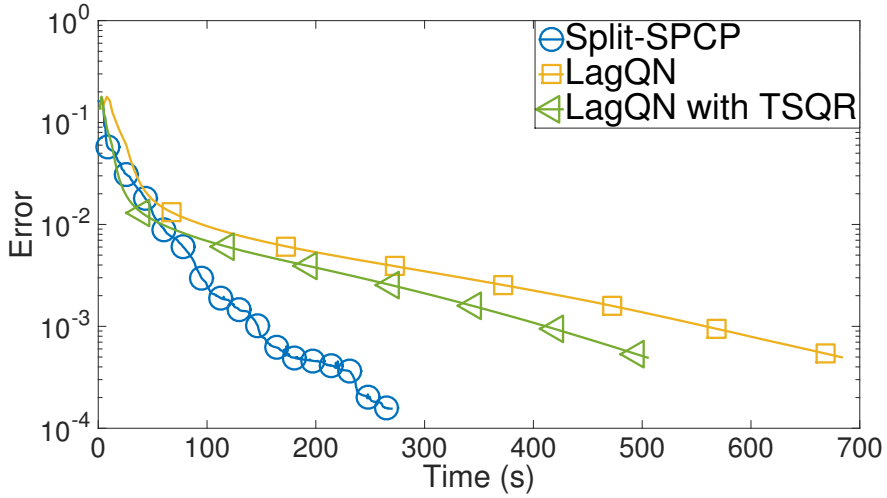

Figure 17: Comparing Split-SPCP with the convex solver LagQN with TSQR and without TSQR. Split-SPCP was initialized with the rank bound $k = 462$, the maximum possible rank of the low rank component. One marker represents 35 iterations.

As with previous tests, a reference solution (L_{ref}, S_{ref}) was found by solving LagQN to high accuracy, and the error is measured as the normalized distance from the reference objective value. The rank of L_{ref} was 423, and Split-SPCP was initialized with the rank bound $k = 462$. Since X had 462 columns, this was the largest possible rank bound, and Split-SPCP still greatly outperformed the convex solvers. Also, both solvers found solutions of similar quality. At an error of 5×10^{-4}, the low rank component found by LagQN had rank 427, and the low rank component found by Split-SPCP had rank 425. Similarly, the sparse component of LagQN had sparsity 58.5%, and the sparse component of Split-SPCP had sparsity 58.3%, while S_{ref} was

58.5% sparse. Non-convex SPCP recovered nearly the same solution as the parallelized convex algorithm at a much smaller cost.

7 Conclusion

In this manuscript, we developed two new methods to accelerate regularized low rank recovery models. Both of these methods offer significant speedup over existing approaches, and both models are amenable to further speedup using parallelized computational architectures.

For our first approach (Split-SPCP), we use the factorization $L = UV^T$ to induce an upper-bound on the rank of L and eliminate the iterative SVDs normally required for nuclear norm minimization. By marginalizing the regularized S variable, we create an objective that is Lipschitz-differentiable when the loss function is strongly convex. Using L-BFGS on the GPU, we showed that our non-convex SPCP model converges much faster than existing methods. We also developed a certificate to determine whether our model had converged to a minimum or to a spurious stationary point, showing empirically and theoretically that local minima do not pose a serious problem for our non-convex model.

Our second approach (TSQR) maintained the convexity of the low rank recovery model and accelerated the iterative SVDs by reducing the algorithm's communication. Combining the tall-skinny QR decomposition with the rSVD algorithm, we solved an inherent communication problem that exists when running the rSVD on a GPU. This reduction in communication accelerated convex SPCP solvers by 40% in our experiments. While the non-convex model outperformed TSQR, it is an important contribution, showing how to better parallelize the original convex formulation on GPU architectures.

Our accelerated solvers provide regularized low rank recovery models with new applications. The non-convex SPCP approach can process video streams in real-time. Both of our models allow for faster decomposition of extremely large data sets, such as fMRI brain scans, which makes the low rank recovery of these data sets feasible.

8 Acknowledgments

We would like to thank Tor Wager and Stephan Geuter from the Neuroscience department of the University of Colorado-Boulder for the fMRI data and advice.

References

Anderson, Michael, Ballard, Grey, Demmel, James, and Keutzer, Kurt. Communication-avoiding QR decomposition for GPUs. In *GPU technology conference*, 2010.

Aravkin, A., Becker, S., Cevher, V., and Olsen, P. A variational approach to stable principal component pursuit. In *Uncertainty in Artificial Intelligence*, Quebec City, 2014a.

Aravkin, Aleksandr, Kumar, Rajiv, Mansour, Hassan, Recht, Ben, and Herrmann, Felix J. Fast methods for denoising matrix completion formulations, with applications to robust seismic data interpolation. *SIAM Journal of Scientific Computing*, 36(5):S237—-S266, 2014b.

Aravkin, A.Y., Burke, J.V., Drusvyatskiy, D., Friedlander, M.P., and Roy, S. Level-set methods for convex optimization. *arXiv preprint*, 2016.

Bach, F., Jenatton, R., Mairal, J., and Obozinski, G. Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.*, 4(1):1–106, January 2012. ISSN 1935-8237. URL http://dx.doi.org/10.1561/2200000015.

Bach, Francis, Mairal, Julien, and Ponce, Jean. Convex sparse matrix factorizations. *Technical report, preprint arXiv:0812.1869*, 2008.

Bauschke, H. H. and Combettes, P. L. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer-Verlag, New York, 2011.

Bogdan, Malgorzata, van den Berg, Ewout, Sabatti, Chiara, Su, Weijie, and Candès, Emmanuel. SLOPE-adaptive variable selection via convex optimization. *The Annals of Applied Statistics*, 9(3):1103–1140, 2015.

Bouwmans, Thierry and Zahzah, El-hadi. Robust PCA via principal component pursuit: A review for a comparative evaluation in video surveillance. *Computer Vision and Image Understanding*, 2014.

Burer, S. and Monteiro, R.D.C. Local minima and convergence in low-rank semidefinite programming. *Math. Prog. (series A)*, 103(3):427–444, 2005.

Candès, E. J., Li, X., Ma, Y., and Wright, J. Robust principal component analysis? *Journal of ACM*, 58 (1):1–37, 2011.

Candès, E. J., Li, X., and Soltanolkotabi, M. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.

Davis, Damek. An $O(n \log(n))$ algorithm for projecting onto the ordered weighted ℓ_1 norm ball. Technical Report CAM 15-32, University of California, Los Angeles, 2015.

Demmel, J., Grigori, L., Hoemmen, M., and Langou, J. Communication-optimal parallel and sequential QR and LU factorizations. *SIAM J. Sci. Comput.*, 34:A206—-A239, 2012.

Fazel, Maryam, Hindi, Haitham, and Boyd, Stephen P. A rank minimization heuristic with application to minimum order system approximation. In *American Control Conference, 2001. Proceedings of the 2001*, volume 6, pp. 4734–4739. IEEE, 2001.

Feia, Yun, Rongb, Guodong, Wanga, Bin, and Wangc, Wenping. Parallel L-BFGS-B algorithm on gpu. *Elsevier*, 40, 2014.

Frank, M. and Wolfe, P. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3 (1):95—110, 1956.

Gunasekar, S., Acharya, A., Gaur, N., and Ghosh., J. Noisy matrix completion using alternating minimization. In *Machine Learning and Knowledge Discovery in Databases*, pp. 194—-209. Springer, 2013.

Halko, N., Martinsson, P. G., and Tropp, J. A. Finding structure with randomness: Stochastic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.

Harchaoui, Z., Juditsky, A., and Nemirovski, A. Conditional gradient algorithms for norm-regularized smooth convex optimization. *Mathematical Programming*, 152(1):75–112, 2015.

Hardt, M. On the provable convergence of alternating minimization for matrix completion. *arXiv preprint*, 2013.

Jaggi, M. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, 2013.

Jaggi, M. and Sulovsk, M. A simple algorithm for nuclear norm regularized problems. In *International Conference on Machine Learning (ICML)*, 2010.

Jain, P., Netrapalli, P., and Sanghavi, S. Low-rank matrix completion using alternating minimization. In *Proceedings of the 45th annual ACM Symposium on the Theory of Computing*, pp. 665—-674. ACM, 2013.

Keshavan, R. *Efficient algorithms for collaborative filtering*. PhD thesis, Stanford University, 2012.

Keshavan, R. and Montanari, A. Regularization for matrix completion. In *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 1503—-1507, 2010.

Keshavan, R., Montanari, A., and Oh, S. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, pp. 952—-960, 2009.

Keshavan, R., Montanari, A., and Oh, S. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56:2980–2998, 2010.

Li, L., Huang, W., Gu, I., and Trian, Q. Statistical modeling of complex backgrounds for foreground object detection. *IEEE Transaction on Image Processing*, 2004.

Lin, Zhouchen, Liu, Risheng, and Su, Zhixun. Linearized alternating direction method with adaptive penalty for low-rank representation. In *Advances in Neural Information Processing Systems*, 2011.

Loader, Clive. *Local Regression and Likelihood*. Springer, 1999.

Martinsson, Per-Gunnar and Voronin, Sergey. A randomized blocked algorithm for efficiently computing rank-revealing factorizations of matrices. *arXiv*, 2015. URL https://arxiv.org/abs/1503.07157.

Mu, Cun, Zhang, Yuqian, Wright, John, and Goldfarb, Donald. Scalable robust matrix recovery: Frank-wolfe meets proximal methods. *SIAM Journal on Scientific Computing*, 2016.

Natarajan, B. K. Sparse approximate solutions to linear systems. *SIAM Journal of Computing*, 24:227—-234, 1995.

Netrapalli, Praneeth, Niranjan, U. N., Sanghavi, Sujay, Anandkumar, Animashree, and Jain, Prateek. Non-convex robust PCA. In *Neural Information Processing Systems*, 2014.

Recht, B., Fazel, M., and Parrilo, P. Guaranteed minimum rank solutions of matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

Rockafellar, R. and Wets, R. *Variational Analysis*. Springer Berlin Heidelberg, 2009.

Rodrígues, P. and Wohlberg, B. Fast principal component pursuit via alternating minimization. In *Proceedings of IEEE international conference on image processing (ICIP)*, pp. 69–73, 2013.

Schmidt, M. minFunc: unconstrained differentiable multivariate optimization in Matlab. 2005. URL http://www.cs.ubc.ca/~schmidtm/Software/minFunc.html.

Shen, Y., Wen, Z., and Zhang, Y. Augmented Lagrangian alternating direction method for matrix separation based on low-rank factorization. *Optimization Methods and Software*, 2012.

Sobral, Andrews and Vacavant, Antoine. A comprehensive review of background subtraction algorithms evaluated with synthetic and real videos. *Computer Vision and Image Understanding*, 122:4–21, 2014.

Srebro, N., Rennie, J., and Jaakkola, T. Maximum-margin matrix factorization. In *NIPS*, 2005.

Tomov, Stanimire, Dongarra, Jack, and Baboulin, Marc. Towards dense linear algebra for hybrid GPU accelerated manycore systems. *Parallel Computing*, 36(5-6):232–240, June 2010. ISSN 0167-8191. doi: 10.1016/j.parco.2009.12.005.

Tu, Stephen, Boczar, Ross, Soltanolkotabi, Mahdi, and Recht, Benjamin. Low-rank solutions of linear matrix equations via procrustes flow. In *International Conference on Machine Learning*, 2016.

Udell, Madeleine, Horn, Corinne, Zadeh, Reza, and Boyd, Stephen. Generalized low rank models. *Foundations and Trends in Machine Learning*, 9(1):1–118, 2015.

Vacavant, Antoine, Chateau, Thierry, Wilhelm, Alexis, and Lequiévre, Laurent. A benchmark dataset for foreground/background extraction. In *ACCV 2012, Workshop: Background Models Challenge*, pp. 291–300, 2012.

Wen, Z., Yin, W., and Zhang, Y. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation*, pp. 1–29, 2010.

Wright, S. J. *Primal-Dual Interior-Point Methods*. SIAM, 1997.

Wright, Y., Ganesh, A., Rao, S., and Ma, Y. Robust principal component analysis: Exact recovery of corrupted low-rank matrices by convex optimization. In *Neural Information Processing Systems (NIPS)*, 2009.

Zeng, Xiangrong and Figueiredo, Mário A. T. The ordered weighted l1 norm: Atomic formulation, dual norm, and projections. *arXiv*, 2014. URL https://arxiv.org/pdf/1409.4271v5.pdf.

Zhou, Tianyi and Tao, Dacheng. Godec: Randomized low-rank & sparse matrix decomposition in noisy case. In *International Conference on Machine Learning (ICML)*, 2011.

Zhou, Zihan, Li, Xiaodong, Wright, John, Candès, Emmanuel J., and Ma, Yi. Stable principal component pursuit. In *IEEE International Symposium on Information Theory (ISIT)*, 2010.

Zou, H. and Hastie, T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society (B)*, 67:301–320, 2005.

Appendix A: Factorization Theorem Details

Both problems share the same set of local minimizers. This was first observed in the case of SDP in Burer & Monteiro (2005) and later generalized in Aravkin et al. (2014b).

Our variant of the theorem:

Theorem 2. *Consider an optimization problem of the following form*

$$\min_{X \succeq 0} \quad f(X), \quad \textit{such that} \quad \operatorname{rank}(X) \leq k \tag{30}$$

where $X \in \mathbb{R}^{n \times n}$ is a positive semidefinite real matrix, and f is a lower semi-continuous (lsc) function mapping to $[-\infty, \infty]$ and has a non-empty domain over the positive semi-definite matrices. Using the change of variable $X = PP^T$, take $P \in \mathbb{R}^{n \times k}$, and consider the problem

$$\min_{P} \quad g(P) \stackrel{def}{=} f(PP^T) \tag{31}$$

Let $\bar{X} = \bar{P}\bar{P}^T$, where $\bar{X}$ is feasible for (30). Then $\bar{X}$ is a local minimizer of (30) if and only if $\bar{P}$ is a local minimizer of (31).

Proof. We follow ideas from both Burer & Monteiro (2005) and Aravkin et al. (2014b). From Lemma 2.1 in Burer & Monteiro (2005), if both P and K are $n \times k$ matrices, then $PP^T = KK^T$ if and only if $P = QK$ for some orthogonal matrix $Q \in \mathbb{R}^{k \times k}$. The objective in (31) depends only on PP^T, so it is clear that P is a local minimizer of (31) if and only if PQ is a local minimizer for all orthogonal Q.

Note that g defined by $P \mapsto f(PP^T)$ is also lsc since it is the composition of f with the continuous function $P \mapsto PP^T$. We require the functions to be lsc so that the notion of "local minimizer" is well-defined.

Suppose $\bar{X}$ is a local minimizer of (30), i.e., for some $\epsilon > 0$ there is no better feasible point, and factor it as $\bar{X} = \bar{P}\bar{P}^T$ for some $n \times k$ matrix $\bar{P}$. Then we claim $\bar{P}$ is a local minimizer of (31). If it is not, then there exists a sequence $P_n \to \bar{P}$ with $g(P_n) < g(\bar{P})$ for all n. By continuity of the map $P \mapsto PP^T$, there is some n large enough such that $X_n \stackrel{\text{def}}{=} P_n P_n^T$ is within ϵ of $\bar{X}$, with $f(X_n) < f(\bar{X})$, contradicting the local optimality of $\bar{X}$.

We prove the other direction using contrapositive. Suppose $\bar{X}$ is not a local minimizer of (30), so there is a sequence $X_n \to \bar{X}$ with $f(X_n) < f(\bar{X})$. Factor each X_n as $X_n = P_n P_n^T$, and observe that it is not true that P_n converges. However, (X_n) converges and hence is bounded, thus (P_n) is bounded as well (for the spectral norm, $\|X_n\| = \|P_n\|^2$, and over the space of matrices, all norms are equivalent). Since P_n are in a finite dimensional space, so the Bolzano-Weierstrass theorem guarantees that there is a sub-sequence of (P_n) that converges. Let the limit of the sub-sequence be $\bar{P} = \lim_{k \to \infty} P_{n_k}$, and note $\bar{P}\bar{P}^T = \bar{X}$ since $X_{n_k} \to \bar{X}$. Then $g(P_{n_k}) = f(X_{n_k}) < f(\bar{X}) = g(\bar{P})$, so $\bar{P}$ is not a local solution. It also follows from the first paragraph of the proof that there can not be another local solution $\tilde{P}$ that also satisfies $\tilde{P}\tilde{P}^T = \bar{X}$. □

Remark 1. *We recover the settings of both Aravkin et al. (2014b, Thm. 4.1) and Burer & Monteiro (2005), since allowing f to be extended valued and lsc encompasses constraints.*

Remark 2. *Aravkin et al. (2014b, Cor. 4.2) can be clarified to state that the two problems listed there are "equivalent" in the sense that they share the same local minimizers (as well as global), using similar continuity arguments on the mapping $\mathcal{R}$ and its adjoint. Furthermore, since the constraints are compact, solutions exist in both formulations.*